

Further Studies On The Use Of Negative Information In Mobile Robot Localization *

Jan Hoffmann, Michael Spranger, Daniel Göhring, Matthias Jüngel, 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—This paper deals with how the absence of an expected sensor reading can be used to improve Markov localization. *Negative information* has not been used for robot localization for various reasons like sensor imperfections, and occlusions that make it hard to determine if a missing sensor reading is really caused by the absence of a feature. We address these difficulties by carefully modeling the robot’s main sensor, its camera. Taking into account the viewing frustum and detected obstacles, the absence of a sensor reading can be associated with the absence of that particular feature. This information can then be integrated into the localization process. We show the positive effect on robot localization in various experiments. (a) In a specific setup, the robot is able to localize using negative information where without it, it is unable to localize. (b) We demonstrate the importance of modeling occlusions and the impact of false negatives on localization. (c) We show the positive impact in a typical run.

Index Terms—Negative Information, Negative Evidence, Mobile Robots, Markov Localization, Monte Carlo Localization, Obstacles Entropy

I. INTRODUCTION

The classic example of negative information was described in the Sherlock Holmes case “Silver Blaze.” In this case, a house has been broken into. Under such circumstances, one would expect the watch-dog to bark. The curious incident of the non-barking of the dog in the nighttime provides Holmes with the information that the dog must know the burglar, allowing him to solve the case.

We apply this idea to mobile robot localization: we allow the robot to draw conclusions from sensor readings that it expected but did not actually make. As Thrun, Burgard, and Fox put it quite graphically, “not seeing the Eiffel Tower in Paris implies that it is unlikely that we are right next to it” [15]. Using a probabilistic localization approach, this information can be treated like an additional sensor. Markov localization, especially Monte Carlo particle filters, have proven to be highly successful in robot localization in various environments such as office environments [1], dynamic environments such as museums [16] and RoboCup [10], and outdoor applications in unstructured environments [12]. [3] contrasts and benchmarks implementations of Markov localization.

Our work is focussed on localization based on camera images taken with a pan-tilt camera of small opening angle of 55° . These images are processed and landmarks are detected. Whenever the robot senses a landmark, the localization estimate is updated using the sensor model. This sensor model is acquired before the actual run. It describes the probability of the measurement z given a state s (position, orientation, etc.) of the robot. Sensor updates only occur when landmarks are detected. In previous approaches, if no landmark is detected, the state estimation is updated using (only) the motion model of the robot. We extend the localization method to specifically incorporate negative information as proposed by [6].

Negative information is defined as the ascertained absence of an expected sensor reading. The term is chosen in accordance with the terms “(false) positive” and “(false) negative” used in statistics.

Negative information constitutes a smaller information gain per update than sensing a landmark since, in general, there are fewer locations from which a landmark is visible (i.e. high information gain when a landmark is detected) than positions from where it is not (i.e. low information gain when no landmark is detected). A landmark is, by definition, something that stands out in an environment. However, as the robot moves about in its environment, negative information can be integrated over time and can yield significantly improved localization performance. This is especially true for cases where the robot cannot focus on landmarks because of its actual task at hand.

In the above definition, it is important to note that the absence of a sensor reading needs to be ascertained. The absence of a reading on a real robot has three main reasons: i. the target may not be within the sensor range, ii. the object may be occluded, and iii. the sensor may be unable to detect the target because of sensor imperfections, imperfect image processing, etc.. Differentiating the first two cases is not a trivial task and requires careful sensor modeling. We address this problem by considering the field of view (viewing frustum) of the robot and by using obstacle detection to estimate occlusions. The third case is modeled probabilistically in the same way as “regular” sensors where the probability

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

Algorithm 1 Iterative Bayesian updating incorporating negative evidence

- 1: $Bel^-(s_t) \leftarrow \int p(s_t|s_{t-1}, u_{t-1})Bel(s_{t-1})ds_{t-1}$
 - 2: **if** (landmark l detected) **then**
 - 3: $Bel(s_t) \leftarrow \eta p(z_t|s_t)Bel^-(s_t)$
 - 4: **else**
 - 5: $Bel(s_t) \leftarrow \eta p(z_{t,l}^*|s_t, r_t, o_t)Bel^-(s_t)$
 - 6: **end if**
-

to make sensor reading z given the current robot state s is considered.

Related Work. Negative information modeling has been applied to object tracking (see [14] for an introduction and [7] for an overview). The event of not detecting an object is treated as evidence that can be used to update its probability density function [8]. In the RoboCup domain, not seeing the ball on the field can be used to delete Monte Carlo particles in that region as long as occlusions are considered [9]. Negative information is also mentioned in the context of simultaneous localization and mapping (SLAM) where it is used to adjust the confidence in landmark candidates [12]. The notion of negative information for robot localization was introduced recently in [6]. The authors give an introduction and give qualitative experimental proof of the benefit of negative information. We continue this work by giving experimental evidence showing that the concept is robust even in critical situations. We will highlight issues encountered and consider performance of the approach as compared to existing ones, giving a more quantitative analysis.

Outline. In section II, we will briefly summarize how negative information can be incorporated into Monte Carlo localization. This is followed by a detailed description of the sensor model used. In section III the positive impact on localization and the approach’s robustness will be demonstrated in a number of experimental benchmarks using a simulation of the Sony Aibo ERS-7 robot.

II. EXPLOITING NEGATIVE INFORMATION

A. Iterative Bayesian Updating

This work is based on Markov localization for mobile robots as described in [1], [15], [13]. The belief state of the robot $Bel(s_t)$ at time t to be in state s_t is determined by *all* previous robot actions u_t and observations z_t . Using Bayes law and the Markov assumption, $Bel(s_t)$ can be written as a function depending only on the previous belief $Bel(s_{t-1})$, the last robot action u_{t-1} , and the current observation z_t :

$$Bel^-(s_t) \leftarrow \int p(s_t|s_{t-1}, u_{t-1})Bel(s_{t-1})ds_{t-1} \quad (1)$$

$$Bel(s_t) \leftarrow \eta p(z_t|s_t)Bel^-(s_t) \quad (2)$$

with normalizing constant η . Equation 1 shows the *a priori* belief $Bel^-(s_t)$, which propagates the previous belief using the motion model of the robot. The measurement is then




Fig. 1. Illustration of the viewing frustum of a Sony Aibo facing a landmark on a RoboCup field.

incorporated into the belief as described in (2) using the sensor model (‘sensor updating’).

In Markov localization, given an initial belief $Bel(s_0)$ at $t = t_0$, the robot first updates its belief using odometry, and then incorporates new sensor information. The belief is updated iteratively in this fashion for every following time step. In the absence of sensor readings, no sensor updating is performed and the belief is updated solely using odometry.

B. The Notion Of Negative Information

Negative information describes the absence of a sensor reading in a situation where a sensor reading is expected given the current position estimate.

To integrate negative information, imagine a binary sensor being added that fires whenever the primary sensor *does not* detect a particular landmark l . Its probability of it firing is given by:

$$p(z_{l,t}^*|s_t) \quad (3)$$

This probability distribution can be refined by taking into account the sensing range r_t of the robot’s sensors and possible occlusions o_t of landmarks. For a mobile robot with a pan-tilt camera, this volume is a function of the current robot state, i.e. its position, orientation, the pan and tilt angles of the camera and the camera’s opening angle.

By o_t we denote a means of detecting whether or not occlusions have occurred. Occlusion can be caused by the geometry of the environment or by other mobile objects and agents in the environment. The former can be determined with the help of a map of the environment. The latter requires the capability to somehow sense, model, and differentiate objects from the environment.

Combining the two yields the probability of not sensing an expected landmark l at time t :

$$p(z_{t,l}^*|s_t, r_t, o_t) \quad (4)$$

Whenever a landmark is not detected, it can be used in the sensor update step of the Iterative Bayesian Updating (see Algorithm 1).




Fig. 2. Occlusions caused by other robots. The robot is abstracted by the gray box. If the robot was to move to the left, it would (partially) occlude the landmark.

C. Sensor Modeling For The Sony Aibo

1) *Viewing Frustum*: The Aibo ERS-7 is a 4-legged robot with a camera mounted in its head. The camera has a horizontal opening angle of 55° and the robot's head has 3 degrees of freedom (neck tilt, head pan, head tilt). We abbreviate gaze direction by $\varphi = (\varphi_{\text{tilt1}}, \varphi_{\text{pan}}, \varphi_{\text{tilt2}})$. The sensing range is illustrated in fig. 1 and is determined by these three angles and the current *robot pose* (position and orientation). Whether or not an object is visible is calculated considering the camera's viewing frustum and the physical dimensions of the object.

2) *Occlusions*: In order to account for occlusions, we opted for an approach that has been used successfully for detecting obstacles, referred to as 'visual sonar' [4], [11]: The camera image is scanned in vertical scan lines and unoccupied space in the plane of the field is detected since it can only be of green or white color (field lines). Scanning for these colors tells the robot where obstacles are and where there is free space. This can be used to determine if the visibility of the landmark is impaired, i.e. if it is occluded by other robot or some other obstacle. More specifically, if the expected landmark lies in an area where the robot has detected free space, the likelihood of the corresponding pose estimate is decreased. If it lies outside of the detected free space, no inference can be made.

In the RoboCup environment, occlusions by the static environment do not occur unless the robot leaves the field, so they can safely be neglected. The only cause of occlusions are other robots on the field. An important performance factor of the algorithm is the ability of the vision system to differentiate between obstacles that can actually cause occlusions and those which cannot. In the RoboCup environment, the field border is usually considered (and also detected as) an obstacle. It does, however, mark the end of the world for the robot and can not occlude anything of interest to the robot and therefore needs to be treated accordingly in the context of negative information.

On the other hand, robots that are detected may not occlude landmarks as is shown in fig. 2. Based on the intercept theorems, occlusion occurs only if $\frac{(h_o - h_c) * d_l}{d_o} - h_c > h_l$. One important result of this is that obstacles (i.e., other robots) farther away than about 1m cannot occlude landmarks on the field. Taking this into account further increases the performance of the proposed approach.

Taking into account the viewing frustum and possible occlusions the sensor model for not perceiving an expected

landmark (equation 4) becomes:

$$p(z_{t,l}^* | s_t, z_{t,\text{obs}}) \quad (5)$$

where $z_{t,\text{obs}}$ describes the current *obstacle percept* and $s_t = (x_t, y_t, \vartheta_t, \varphi_t)$ the robot state consisting of the *robot pose* (position x_t, y_t , and orientation ϑ_t) and the current gaze direction φ_t .

D. Monte Carlo Localization, Implementation

This work is based on the Monte Carlo localization described in [13], which also serves as a base line implementation. Sensor updating was extended to account for field of view (FOV) and occlusion as described. This also requires sensor updating to be triggered by new camera images regardless of whether or not there was a percept. Before re-sampling, the weight of an individual particle is calculated as follows: Of all landmarks L , the subset of landmarks L' is detected, the subset L^* is expected but not detected, and lastly the subset L° is not detected but was also not expected: $L = L' \cup L^* \cup L^\circ$ and $L^* \cap L' = \emptyset$. The probability of a particle p_i is calculated by multiplying all the likelihoods of all gathered evidences:

$$p_i = \underbrace{\prod_{l \in L'} s_l(\alpha_{\text{measd}}, \alpha_{\text{expd}})}_{\text{detected}} \cdot \underbrace{\prod_{l \in L^*} s_l^*(\varphi, z_{\text{obs}}, \alpha_{\text{expd}})}_{\text{expected and not detected}} \quad (6)$$

The function s_l is an approximation of the sensor model and returns the likelihood of sensing the landmark l at angle α_{measd} for a particle p_i that expects this landmark to be at α_{expd} . Function s_l^* models the probability of not sensing the expected landmark $l \in L^*$ given the current sensing range as determined by gaze direction φ , the robot pose associated with p_i , and the obstacle percept z_{obs} .

III. EXPERIMENTAL RESULTS

The RoboCup Sony 4-Legged League serves as a test bed for our work. In the 4-Legged League, teams of 4 Sony Aibo ERS-7 robots play soccer against each other in a color coded environment (see the official RoboCup web site for details: www.robocup.org). Colored beacons (4 uniquely color coded beacons plus a blue and a yellow goal) and the field lines (similar to the real soccer field lines) serve the robots for localization. In our experiment, unless otherwise stated, only landmarks were used for localization to emphasize the effect of using negative information.

Two quantities can be used when a landmark is seen: its size in the camera image can be used to estimate the distance to the landmark d_l and the relative angle to the landmark (bearing α_l) can be calculated from its position within the image. In practice we only use the bearing because the distance measurement is error prone. Note that this differs from triangulation where distances are used.

In each of the following experiments, the localization module starts with a uniform (random) particle distribution. As the robot receives sensor measurements, the progression of the distribution over time is monitored.

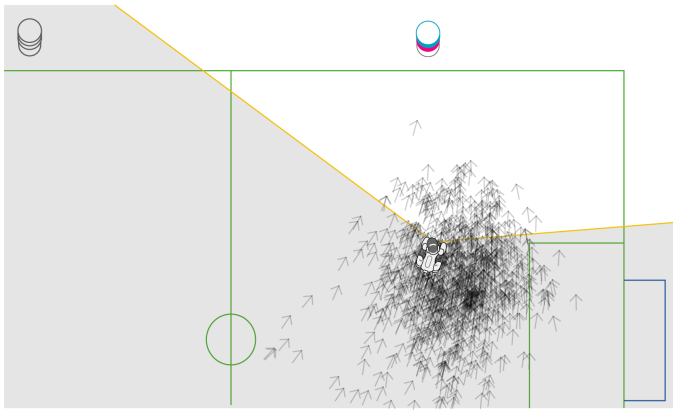


Fig. 3. *Experiment A*: Robot (small white thing in the middle) trying to localize when standing in front of the landmark. The area accessible to the camera when scanning is highlighted. Note that the landmark to the left and the goal on the right are just outside of this area. The particle distribution after 5 seconds using negative information is shown by the arrows.

We use the expected entropy H as an information theoretical quality measure of the position estimate $Bel(s_t)$ [2]:

$$H_p(s_t) = - \sum_{s_t} Bel(s_t) \log(Bel(s_t))$$

The sum runs over all possible states. The entropy of the particle distribution becomes zero if the robot is perfectly localized in one position, maximal values of H mean that $Bel(s_t)$ is uniformly distributed.

In addition to entropy, we consider the distance error of the localization w.r.t. the robot's real position (which is known in simulation). We define the error Δr as the average distance of particles to the actual robot position \vec{x}_{true} :

$$\Delta r = \frac{1}{N} \sum_{i=1}^N |\vec{x}_i - \vec{x}_{true}|$$

where \vec{x}_i is the position of particle i and N is the number of particles. Similarly, entropy and accuracy of the robots orientation is calculated.

The experiments were conducted in simulation using log files to ensure reproducible results and identical sensor input when benchmarking the approach, and to allow for a greater particle number, which results in a smoother representation of the probability distribution.

A. Experiment A: One Landmark

In this experiment, the robot is standing in front of a single landmark and performs a scanning motion with its head. This scan covers $90^\circ + 55^\circ = 145^\circ$ in front of the robot (fig. 3). Within this area there is no other landmark. The goal of this experiment is to show that even with very few sensor readings, localization is possible when negative information is also taken into account. Fig. 4 shows the localization error over time, the number of particles that were updated using negative information and a schematic illustrating the robot's positive




Fig. 4. *Experiment A*: The distance error of the localization over time. The dotted line is without negative information, the solid blue line marked with the star is with negative information integrated. The green bars represent the number of particles, which are updated using negative information (those which expected to see a landmark but did not) summed over all landmarks. Below: The diagram indicates if a landmark (LM) or goal (G) was seen (“|”), not seen (“-”), or expected but not seen for any of the particles (“x”).

and negative percepts. This first experiment reproduces the results presented in [6].

Not using negative information. The experiment starts with a uniform particle distribution, which equals to maximum entropy. When the landmark comes into view, a decrease in localization error can be observed. This information gain is due to the robot being able to now infer its relative orientation w.r.t. the landmark. Since there are no further constraints on the robot's position, the entropy remains at a relatively high level and the position error does not decrease. In other words, although there is a drop in entropy, the position estimate is still highly uncertain.

Incorporating negative information. When using negative information, the particle distribution converges towards the actual robot pose. If you look at fig. 3 closely, you will notice that the range of the scan covers quite a lot of the field and that it is bordering on a goal and a beacon. This has an effect on how much the particle distribution converges. In the case described here, almost a maximum of negative information is incorporated. The maximum would be reached if the scan range could not be increased any further without including the landmarks. In this case, the localization result would be just as good as actually seeing the landmarks. If the range of the scan was smaller or the robot would stand closer to the landmark, less negative information would be gathered,

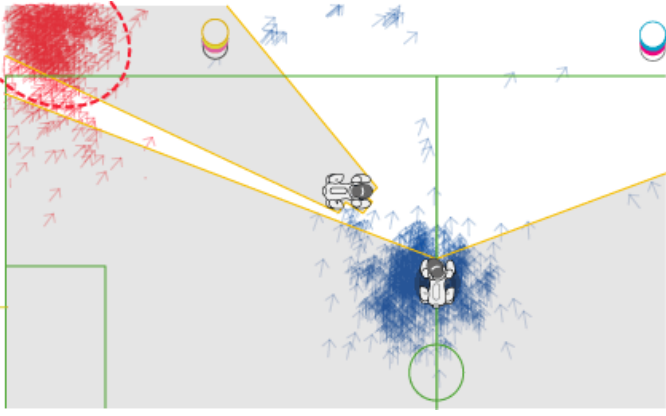


Fig. 5. *Experiment B*: In this situation, the landmark on the left is occluded by another robot. The red rectangle marks the actual position of the robot on the field close to the center circle facing outward. The blue particle distribution shows the localization using negative information and taking occlusion into account. The red particle distribution (also marked by the dashed line) in the top left corner shows the effect of false negatives: if obstacles are not treated properly, false negatives incorporated into the probability distribution results in convergence in the wrong area of the field.

and the distribution would converge only to an area around a segment of a circle around the landmark.

Note that the number of particles that are updated using negative information decreases as the localization improves.

B. Experiment B: Two Landmarks, One Occluded

The goal of this experiment was to prove that obstacles are modeled correctly. The robot is placed on the center line and it performs a scan as shown in fig. 5. From where it is standing, two landmarks are within its sensing range. However, one of the landmarks is occluded by another robot.

As in the previous experiment, the standard approach is unable to localize. When negative information is incorporated taking into account occlusions, the robot is able to localize quite accurately.

The figure also shows what happens if obstacles are not modeled correctly. Not modeling obstacles causes *false negatives*, i.e. the robot fails to see a landmark due to occlusions and wrongly assumes that it is absent. In these cases, the particle distribution converges to a completely wrong robot pose on the corner of the field. (This position is, of course, compatible with seeing a landmark at the right end of the scan an not seeing anything for the rest of the scan.)

C. Experiment C: Moving Robot

In this experiment, the robot walks on field following the ball. The aim of it is to show the performance of the approach in an actual application. Fig. 6 shows the localization error and the particle distribution entropy in the first 3 seconds of the run.

In all four curves, the appearance of landmarks in the camera images leads to an improvement in localization. When only percepts are used, distinct steps can be seen in the respective curves. As long as no new evidence is gathered, the level of uncertainty stays the same.

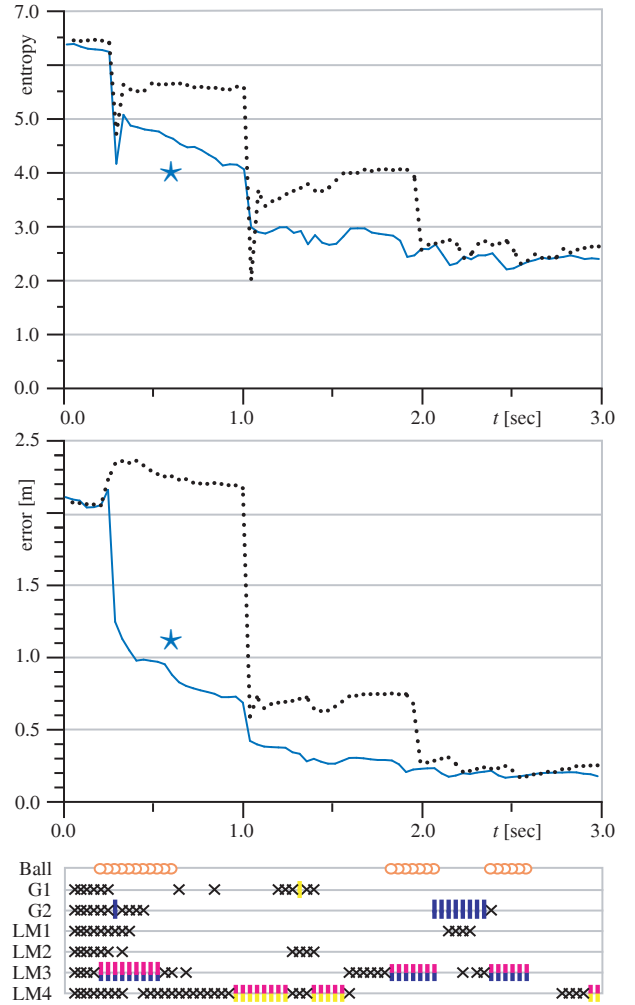


Fig. 6. *Experiment C*: Entropy of the particle distribution and localization error of a robot walking on the field chasing a ball. The solid, blue lines marked with the star represent the localization with negative information used, dotted lines without negative information. It shows that negative information is able to “fill in the blanks” before the next landmark is actually perceived.

Incorporating negative information leads to a smoother decrease in uncertainty and better localization in the case of limited percepts. After some time, though, the robot has seen three landmarks and the quality of the localization reaches similar levels in both cases.

One important conclusion from this is that negative info can help fill in the blanks in situations where there is incomplete sensor input. The quality of the resulting localization is limited by the best possible localization using all percepts potentially available at a given position. In other words: negative and positive evidence are two sides of the same coin; a well localized robot cannot further improve its localization by negative evidence because there will be no negative evidence (cf. number of particles updated using negative information in fig. 4).

D. Kidnapped Robot

The kidnapped robot problem is a commonly used benchmark for the flexibility and robustness of localization algorithms [3]: a localized robot is displaced and the time for it to recover is measured.

As indicated in [6], the increased responsiveness of the localization when using negative information has a positive impact in the kidnapped robot benchmark. One reason for this is the additional information that is used when the robot re-localizes, resulting in faster convergence of the particle distribution as shown in the above experiments. But not only does the distribution converge more quickly, it also *diverges* quicker in the absence of sensor readings that would confirm the robot's past position. The distribution therefore better resembles the actual situation of being lost and offers a much better starting point for subsequent re-localization. This is particularly helpful in RoboCup game situations where the robot often gets pushed by other robots. Unless collisions are explicitly modeled [5], these relatively small displacements may go unnoticed by the localization. Using negative information, the particle distribution diverges quickly to nearby positions, which often quite accurately models what has actually happened to the robot.

Run time performance. No explicit run time measurements have been made so far since the code is not yet optimized to run on the Aibo. Qualitatively speaking, the calculations necessary to decide if negative evidence is present is slightly more complex than the calculations used to calculate the similarity of a percept. These calculations have to be made for all 6 landmarks (compared to 1-2 for seen landmarks). In the worst case, this means that the weighting of particles takes 6 times as long.

IV. CONCLUSION

Negative evidence can be used to improve Markov localization of a mobile robot. To achieve this, the robot's sensor needs to be modeled carefully. In our case, this means taking into account the robot's camera's viewing frustum and obstacles in its environment. In experiments we are able to show the strength of this approach. It allows the robot to localize in situations where without the use of negative information it could not. We also show how obstacle modeling helps to avoid "false negatives". In actual application scenarios, negative information is able to fill in the gaps when not enough landmarks are seen. It is interesting to note that in this context, it may make sense for the robot to direct its gaze towards areas where it doesn't actually see a landmark but can rule out hypothesis that would expect a landmark in that area.

Future work. We will work on ways to improve run time performance by not considering the potential negative evidence of all landmarks but only of those relevant. This work will be part of an active vision system for the robot where early tests have shown that modeling negative evidence is a good way to decide if sensing actions are successful and thus guide attention.

ACKNOWLEDGMENTS

Program code used was developed by the GermanTeam, a joint effort of the Humboldt University of Berlin, University of Bremen, University of Dortmund, and the Technical University of Darmstadt. Source code is available for download at <http://www.germanteam.org>. Freerk Stulp pointed out "Silver Blaze" to us.

REFERENCES

- [1] D. Fox, W. Burgard, F. Dellart, and S. Thrun. Monte carlo localization: Efficient position estimation for mobile robots. In *Proc. of AAAI*, 1999.
- [2] D. Fox, W. Burgard, and S. Thrun. Active markov localization for mobile robots. In *Robotics and Autonomous Systems*, 1998.
- [3] J.-S. Gutmann and D. Fox. An experimental comparison of localization methods continued. *Proceedings of the 2002 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2002.
- [4] J. Hoffmann, M. Jüngel, and M. Löttsch. A vision based system for goal-directed obstacle avoidance. In D. Nardi, M. Riedmiller, C. Sammut, and J. S.-V. (Eds.), editors, *8th International Workshop on RoboCup 2004 (Robot World Cup Soccer Games and Conferences)*, volume 3276 of *Lecture Notes in Artificial Intelligence*, pages 418–425. Springer, 2005.
- [5] J. Hoffmann, M. Spranger, D. Göhring, and M. Jüngel. Exploiting the unexpected: Negative evidence modeling and proprioceptive motion modeling for improved markov localization. In *9th International Workshop on RoboCup 2005 (Robot World Cup Soccer Games and Conferences)*. Springer, 2006. to appear.
- [6] J. Hoffmann, M. Spranger, D. Göhring, and M. Jüngel. Making use of what you don't see: Negative information in markov localization. In *Proceedings of the IEEE/RSJ International Conference of Intelligent Robots and Systems (IROS) 2005*, 2006. to appear.
- [7] W. Koch. On negative information in tracking and sensor data fusion. In *Proceedings of the Seventh International Conference on Information Fusion*, pages 91–98, 2004.
- [8] W. Koch. Utilizing negative information to track ground vehicles through move-stop-move cycles. In *Proceedings of the SPIE*, volume 5429, pages 273–283, 2004.
- [9] C. Kwok and D. Fox. Map-based multiple model tracking of a moving object. In *8th International Workshop on RoboCup 2004 (Robot World Cup Soccer Games and Conferences)*, *Lecture Notes in Artificial Intelligence*. Springer, 2005. to appear.
- [10] 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.
- [11] S. Lenser and M. Veloso. Visual sonar: Fast obstacle avoidance using monocular vision. In *Proceedings of IROS'03*, 2003.
- [12] M. Montemerlo and S. Thrun. Simultaneous localization and mapping with unknown data association using fastslam. 2003.
- [13] T. Röfer and M. Jüngel. Vision-based fast and reactive monte-carlo localization. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA-2003), Taipei, Taiwan*, pages 856–861, 2003.
- [14] S. Särkkä, T. Tamminen, A. Vehtari, and J. Lampinen. Probabilistic Methods in Multiple Target Tracking, Research Report B36. Technical report, Laboratory of Computational Engineering Helsinki University of Technology, 2004.
- [15] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*, page 231. MIT Press, 2005.
- [16] S. Thrun, D. Fox, and W. Burgard. Monte carlo localization with mixture proposal distribution. In *Proc. of the National Conference on Artificial Intelligence*, pages 859–865, 2000.