

Real-Time Active Vision by Entropy Minimization Applied to Localization

Stefan Czarnetzki, Sören Kerner, and Michael Kruse

Robotics Research Institute
Section Information Technology
TU Dortmund University
44221 Dortmund, Germany

Abstract. This paper presents an active vision approach to enhance mobile robot localization. A particle filter localization is extended with a module to find active vision decisions that are optimal based on the current localization and its uncertainty. Optimality is expressed as a criterion of entropy minimization. Further approximations are introduced to enable real-time computation. Both the usefulness of the presented approach in a RoboCup scenario and the performance and quality of the approximations are evaluated in different static and dynamic situations.

1 Introduction

Knowing their own position relative to the environment and certain points of interest is essential for mobile robots to autonomously achieve their tasks. Approaches focusing on localization exclusively are called *passive localization*, i.e. they address the position estimation based on an incoming stream of sensor data that can not be influenced. Solutions based on Kalman filters, particle filters and combinations and extensions thereof are numerous and common in this field. A comprehensive overview is given in [1]. *Active localization* instead implies the interaction of the localization process with the control of the robot. Examples for control or influence of the whole navigation process are coastal navigation algorithms where the localization uncertainty is anticipated and taken into account for path planning [2]. The more complex the robot's task is, however, the more problematic to include localization quality as an additional criterion into the general planning processes for achieving the given task. A common trade-off is to allow (partial) control over the information gathering process, e.g. viewpoint and orientation of directed sensing devices, while leaving the navigation decisions unaffected. In case of imaging sensors this belongs to the field of *active vision*.

Active vision in general refers to actuated camera systems with a task-oriented control of the camera movement. This might be staying on target or automatically choosing regions of interest for tracking or surveillance applications [3] or fixating on particular objects in remote collaboration scenarios [4].

Active vision for localization tasks means to move the camera with regard to the expected information gain. Therefore it is necessary to specify such information gain, to infer which observations would improve knowledge the most,

and estimate the result of different actions based on possibly multi-modal localization belief states. Previous work exists for different fields of application. The approaches depend on the range of possible actions, the types of observations and the representations of the localization belief [5,6,7].

This paper presents the application of an approach similar to [8] to a humanoid robot with directed vision in the context of a known environment with mainly ambiguous landmarks. First, the underlying localization is described briefly together with concepts for the measurement and estimation of localization quality to infer expected information gains. Modeling of the environment and sensors is necessary to reason about expected results of selected actions. Furthermore the algorithm to find optimal active vision decisions is described in detail including necessary discretization and adaptations for real-time processing on platforms with restrictive resources. Our approach differs from [8] in the underlying localization approach of a particle filter instead of a grid localization and the use of directed vision instead of an array of sonar sensors, which therefore results in different adaptations and approximations for an efficient implementation. Finally the performance both in static and dynamic situations is evaluated and a conclusion is given.

2 Modeling

The computation of an optimal active vision decision depends on the localization task and its belief representation, the actions to choose from and the estimation of their outcomes. The latter involves using knowledge about environment and sensors to first predict observations given a localization belief state and then predict their influence on said belief state. Those predictions obviously depend on the underlying localization approach and models of environment and sensors. This will be described briefly in the following sections.

2.1 Localization

The task of (passive) localization consists of the estimation of the robot's location relative to a known map given a stream of sensor and control information, y_1, \dots, y_t and u_1, \dots, u_t , respectively. Global in contrast to local localization assumes no a-priori knowledge about the initial position x_0 . The current belief state of a robot's localization is therefore given by

$$\text{bel}(x_t) = p(x_t | u_{1:t}, y_{1:t}) \quad (1)$$

and the posterior prior to incorporating the latest observation y_t , the so called prediction, is

$$\overline{\text{bel}}(x_t) = p(x_t | u_{1:t}, y_{1:t-1}). \quad (2)$$

Having to work on the whole collection of sensor and control data from all previous time steps results in complexity and needed memory increasing with time. The Markov assumption of a complete representation of the state x_{t-1} at

the last time step $t - 1$ allows the estimation of the current state x_t based on this last state and the current control input u_t only. The recursive state estimation of the Bayes filter is then given by

$$\overline{bel}(x_t) = \int p(x_t|u_t, x_{t-1}) bel(x_{t-1}) dx_{t-1} \quad (3)$$

$$bel(x_t) = \frac{p(y_t|x_t) \overline{bel}(x_t)}{p(y_{t+1}|u_t)} = \eta p(y_t|x_t) \overline{bel}(x_t) \quad (4)$$

with η normalizing the belief to add up to 1. The state transition $p(x_t|u_t, x_{t-1})$ and the measurement probability $p(y_t|x_t)$ depend on the given task, in case of localization on the robot's odometry and kinematic model and its sensor model. The actual implementation of equations 3 and 4 depends on the representation of $bel(x_t)$.

In case of particle filters the belief is given as a set of particles or hypotheses $X_t = x_t^1, x_t^2, \dots, x_t^n$ representing the distribution

$$x_t^i \propto p(x_t|u_{1:t}, y_{1:t}). \quad (5)$$

2.2 Entropy of Belief States

To evaluate effects of possible actions it is necessary to specify criteria for the usefulness of their result. With respect to localization this refers to the quality of the belief state. Without knowledge about the robot's real position it is only possible to judge the belief state's uncertainty, i.e. the covariance matrix in case of an unimodal Gaussian distribution for Kalman filters. A common measure for this is the entropy.

Entropy $H(X)$ is defined as the expected information content

$$H(X) = - \int p(x) \log p(x) dx. \quad (6)$$

with $p(x) \log p(x) = 0$ for $p(x) = 0$. This allows to compute a single scalar value as a measure for the uncertainty of a given probability distribution. Obviously for a distribution which is non-zero only at a single point, i.e. a gaussian with zero variance or a Dirac delta function, the entropy $H(X) = 0$. For a uniform distribution on the other hand the entropy is maximal.

In case of a probability distribution given as a particle set the analogy of this concept to thermodynamic entropy is obvious. For computing it using equation 6 the value of $p(x)$ can be extracted from the local particle density relative to the overall number of particles.

3 Active Vision Decision Finding

Choosing an action in active vision means to choose where to point the camera. This is comparatively easy as long as the robot's (or the camera's) position is

known. A naive approach would be to choose a feature from the map, e.g. the most descriptive or the closest, and to point the camera in the direction of the expected position. While this might still work as long as the belief distribution is concentrated close to the true robot pose, it becomes less effective once the distribution is less certain or even multi-modal.

The following section describes the computation of an optimal active vision decision under uncertain belief distributions.

3.1 The Optimal Action

To choose an optimal action for a given localization belief state, it is necessary to evaluate the profit of every possible candidate. For every action the probable observations need to be computed, their influence on the belief state inferred and the result evaluated. The presented approach is based on and follows the outline of [8].

In general possible actions might be simple motor commands to position sensors, short motion sequences or complex actions like navigation commands to move the robot to certain positions. In the following every action u_i from the set of possible actions $U = u_1, u_2, \dots, u_m$ will be considered atomic and consequences will only be considered for the complete execution of an action. The notation of u is identical to the control data in section 2.1 since the following holds for robot actions in general. The specific application to active vision is referred to in sections 3.2 to 3.4. The presented approach always chooses the best action based on the immediate benefit and can thus be characterized as greedy.

Section 2.2 illustrates the relationship between entropy and uncertainty. The optimal action u_{opt} for a given belief X_t is the one minimizing the entropy $H(X_{t+1}|u_{opt}, y_{t+1})$ of the belief X_{t+1} after execution of u_{opt} and incorporation of the observation y_{t+1} resulting from this action (see equation 7).

$$u_{opt} = \operatorname{argmin}_{u \in U} H(X_{t+1}|u, y_{t+1}) \quad (7)$$

The terminology here is that only one observation is made at any time, but this single observation might include measurements of a various number of landmarks or features. This observation y_{t+1} , however, can only be predicted with certainty if the true position is known. An uncertain position given as a probability distribution only allows a decision based on the expected entropy of the next belief, as shown in equation 8.

$$\hat{u}_{opt} = \operatorname{argmin}_{u \in U} E[H(X_{t+1}|u)] \quad (8)$$

Computing this expected entropy necessitates the simulation of all possible observations y_{t+1} for a given action. Depending on the complexity of the observation's representation and the sensing process, generating the observations can be based on recorded data as in [5] or [6]. For laser or sonar sensors this might be done using the known environment model and ray-casting operations [7] or a simple visibility simulation for cameras.

Let $p(x)$ from equation 6 be represented by $bel(x)$ and let $p(y_{t+1}|u)$ be the probability of observing y_{t+1} after executing u . Then the expected entropy can be calculated as follows using equations 6 and 4:

$$\begin{aligned} E[H(X_{t+1}|u)] \\ = \int p(y_{t+1}|u) H(X_{t+1}|u, y_{t+1}) dy_{t+1} \end{aligned} \quad (9)$$

$$= - \iint p(y_{t+1}|u) bel(x_{t+1}) \log bel(x_{t+1}) dx_{t+1} dy_{t+1} \quad (10)$$

$$\begin{aligned} &= - \iint p(y_{t+1}|u) \frac{p(y_{t+1}|x_{t+1}) \overline{bel}(x_{t+1})}{p(y_{t+1}|u)} \\ &\quad \log \frac{p(y_{t+1}|x_{t+1}) \overline{bel}(x_{t+1})}{p(y_{t+1}|u)} dx_{t+1} dy_{t+1} \end{aligned} \quad (11)$$

$$= - \iint p(y_{t+1}|x_{t+1}) \overline{bel}(x_{t+1}) \log \frac{p(y_{t+1}|x_{t+1}) \overline{bel}(x_{t+1})}{p(y_{t+1}|u)} dx_{t+1} dy_{t+1}. \quad (12)$$

By inclusion of $\overline{bel}(x_{t+1})$ possible movements of the robot are already considered. If u includes motion controls then choosing those according to equation 12 results in navigation optimizing the robot's localization. If u includes only such commands without influence on the state x_t then those can be controlled to optimize localization without interfering with the robot's task-oriented navigation. When u has no influence on x_{t+1} as in active vision then $\overline{bel}(x_{t+1})$ is the same for all u and needs to be computed only once. In the case that the difference between x_t and x_{t+1} can be neglected, i.e. the robot's motion during this time is small, further simplification is possible resulting in equation 13 omitting calculation of the process model.

$$E[H(X_{t+1}|u)] = - \iint p(y_{t+1}|x_t) bel(x_t) \log \frac{p(y_{t+1}|x_t) bel(x_t)}{p(y_{t+1}|u)} dx_t dy_{t+1} \quad (13)$$

Actions may also vary in necessary effort so that their costs need to be considered. This is done by weighted addition to the term to be minimized in equation 8.

3.2 Discretization

The integrals in equations 12 or 13 can not be computed in closed form for any non-trivial applications. Discretization is necessary at certain points to allow the computation of those terms.

A first step is obvious for the application of a particle filter. Let $\bar{x}_{t+1}^i \in \overline{X}_{t+1}$ be a particle from the set X_t after applying u , then equation 12 becomes

$$\begin{aligned} E[H(X_{t+1}|u)] \\ = - \int \sum_{i=1}^n p(y_{t+1}|\bar{x}_{t+1}^i) p(\bar{x}_{t+1}^i|u) \log \frac{p(y_{t+1}|\bar{x}_{t+1}^i) p(\bar{x}_{t+1}^i|u)}{p(y_{t+1}|u)} dy_{t+1}. \end{aligned} \quad (14)$$

In this case $p(\bar{x}_{t+1}^i|u)$ can not be inferred from a single particle but from the local particle density around it. A working particle localization normally tends to result in few areas with high probability after a certain time. A further approximation might concentrate on those particle clusters and use average positions computed from these subsets instead of all particles separately which lowers the computational cost at least by an order of magnitude and provides the needed local particle density at the same time.

Discretization of the observations can not be motivated by the particle filter concept but depends on the sensor type and the possible observations to be made. For vision based perception those observations are commonly recognized landmarks. A complete enumeration of all possible observations would need to account for all landmarks with all possible distances and bearings which is clearly not practical. To avoid the generation of many observations whose measurement probability is near zero for the current belief an alternative is the simulation of observations for all position hypotheses given by the particle set or the chosen position discretization. This generates a number of observations in the order of the number of landmarks times the amount of position hypotheses. This discretization step results in

$$E[H(X_{t+1}|u)] = - \sum_{j=1}^m \sum_{i=1}^n p(y_{t+1}^j|\bar{x}_{t+1}^i) p(\bar{x}_{t+1}^i|u) \log \frac{p(y_{t+1}^j|\bar{x}_{t+1}^i) p(\bar{x}_{t+1}^i|u)}{p(y_{t+1}^j|u)} \quad (15)$$

or

$$E[H(X_{t+1}|u)] = - \sum_{j=1}^m \sum_{i=1}^n p(y_{t+1}^j|x_t^i) p(x_t^i) \log \frac{p(y_{t+1}^j|x_t^i) p(x_t^i)}{p(y_{t+1}^j|u)} \quad (16)$$

for the case that the robot's motion can be neglected.

Finally the actions to be evaluated in equation 8 need to be enumerated. In case of active vision these are motor commands for all possible pan and tilt angles to control the camera's field of view. Since it is not possible to process this search space in any more efficient way than to do a complete evaluation the chosen discretization is essential for real-time computation.

3.3 Decision Computation

Based on the considerations of the previous sections it is possible to compute the optimal action according to equation 8. In algorithm 1 the belief $bel(x_t)$ is given as a set of particles X_t and the optimal action \hat{u}_{opt} out of the set of possible actions U is calculated under the assumption of negligible robot motion according to equation 16.

Note that the expression of $\|\mathcal{X}_t^i\|/\|X_t\|$ for the term $p(\tilde{x}_t^i)$ is only valid for clusters representing equally sized volumes of state space, e.g. in case of a regular grid or similar clustering techniques, and when all particles are represented exactly once. The probability $p(y_{t+1}^k|\tilde{x}_t^i)$ in line 14 is the measurement probability in section 2.1 and $p(y_{t+1}^k|u)$ from line 15 is equivalent to the normalizing factor η in equation 4 and can be determined by

Algorithm 1. Active localization by entropy minimization

```

Input: particle set  $X_t$ , set of possible actions  $U$ 
1  $H_{min} = \infty$ 
2 Generate a set of clusters  $\mathcal{X}_t$  from  $X_t$ 
3 foreach  $\mathcal{X}_t^i \in \mathcal{X}_t$  do
4     Calculate average  $\tilde{x}_t^i$  of all  $x_t^j \in \mathcal{X}_t^i$ 
5 end
6 foreach  $u \in U$  do
7     foreach  $\mathcal{X}_t^i \in \mathcal{X}_t$  do
8         Generate observation  $y_{t+1}$  based on  $\tilde{x}_t^i$  and  $u$ 
9          $Y_u = Y_u \cup y_{t+1}$ 
10    end
11     $h = 0$ 
12    foreach  $y_{t+1}^k \in Y_u$  do
13        foreach  $\mathcal{X}_t^i \in \mathcal{X}_t$  do
14             $g = p(y_{t+1}^k | \tilde{x}_t^i) \cdot |\mathcal{X}_t^i| / \|X_t\|$ 
15             $h = h - g \log(g/p(y_{t+1}^k | u))$ 
16        end
17    end
18    if  $h < H_{min}$  then
19         $H_{min} = h$ 
20         $u_{opt} = u$ 
21    end
22 end
23 return  $u_{opt}$ 

```

$$p(y_{t+1}^k | u) = \sum_{i=1}^n p(y_{t+1}^k | \tilde{x}_t^i) bel(\tilde{x}_t^i). \quad (17)$$

This implies the advantage of precomputing all probabilities and their sum before computing the entropy, which was omitted here for clarity.

3.4 Adaptations for Real-Time Processing

The use of clusters instead of separate particles both for the generation of observations and the evaluation of their effect as introduced in the previous section already represents a first approximation. While this allows a computation an order of magnitude more efficient than without this approximation, several other adaptations and simplifications are possible to allow the application on robotic platforms with limited processing power such as the humanoid robot Nao (see section 4).

As long as relevant observations are only caused by static elements of the environment it is possible to pre-compute them to be stored in a lookup table for all robot states and relevant control commands. For 2D localization and pan and tilt controls for a camera this table has a dimension of 5 assuming an approximately constant height of the camera. Of course the limited resolution

in each dimension and the constant height approximation cause further loss of precision which might influence the result. This will be evaluated in section 4.

Similarly the measurement probability might also be pre-computed and stored in a lookup table. Additionally to the 5 dimensions mentioned above this table would need to consider the distance and bearing of the observation and in case of certain features like line crossings on a soccer field the feature's orientation is relevant, too. The resulting 8 dimensions would cause the table to exceed several 100 MB for useful resolutions though, so for embedded platforms with limited memory this is no practical solution.

As can be seen in figure 1 for a common situation the localization belief state often focuses on few areas of high probability. The clustering step in line 2 of algorithm 1 already causes the areas without particles to be neglected. This is valid since those areas correspond to quasi-zero probabilities. Omitting further areas with low but non-zero probability represents an approximation to equation 16 but considerably decreases the computational costs.



Fig. 1. Identification of regions with high particle density

4 Evaluation

The usefulness of this active vision approach to aid localization and the validity of the approximations proposed in section 3.4 are evaluated in several experiments presented in this section. The system used for those experiments is the humanoid robot Nao by Aldebaran Robotics which is used in the RoboCup Standard Platform League. This robot contains a x86 AMD Geode 500 MHz CPU and 256 MB SDRAM. The runtime necessary for algorithm 1 with and without further approximations is presented in table 1. The significant range in those measurements is a direct result of the dependency of computational complexity on localization uncertainty. If a localization belief is less certain more positions and consequently more possible observations need to be evaluated.

The following experiments are performed in a 3D simulation for easy access to ground truth information of the robot position and the repeatability of experiments based on the exact same conditions. The environment is a robot soccer

Table 1. Runtime of active vision module with and without approximations

	Maximum	Average	Minimum
Without Approximations	331.2 ms	38.5 ms	11.5 ms
Pre-calculated Observations	294.3 ms	35.6 ms	9.3 ms
Selective Computing	45.8 ms	4.2 ms	1.2 ms

field as used in Standard Platform League competitions. All processes running on the real robot are also executed in simulation including motion, image processing and cognition. Since all cognition processes run at 15 Hz the algorithm with pre-calculated observations and selective computing is suitable even for platforms with restricted resources like the Nao. Active vision is implemented to decide the target viewing direction, move the head and stare in that direction for one second, then continue to choose the next direction.

4.1 Static Situations

The discretization of possible actions used in all following experiments is 7 different pan angles and 2 tilt angles, one to point the camera at features close to the robot and one for far away observations. The choice is based on the camera's opening angles and slightly overlapping fields of view. This represents a minimal number of distinct actions. A more detailed resolution could provide better decisions, but this would imply higher computational costs.

The first experimental setup shows the validity of the active vision decision process and the applied approximations. Situations are generated which allow decisions that are easy to comprehend. First a robot is placed on the field near a penalty kick mark looking at the center circle until this repetitive exclusive perception generates two symmetrical clusters of position hypotheses. This situation might occur for a robot staring at a ball lying at the kick-off position. Figure 2(a) shows the expected entropy for different action choices with and without approximations. Figure 2(b) illustrates the average position error of all particles after executing each action and processing the resulting observations. The quantitative correlation of the values is clearly visible.

A similar situation occurs frequently in soccer games: When penalized players come back into the game they are placed on the side line in the middle of the field. The side is not known to the robot in advance. In the given particle filter localization this penalty information is handled by placing 40% of the particles near each re-entry spot and distributing the rest randomly to cope with wrong game controller handling. The results for this setup are shown in figure 3.

Additional experiments not described here in detail are set up to evaluate the localization performance starting from total uncertainty in different situations focusing on a comparison between active vision guided localization and the simple approach to move the head from side to side to continuously scan the environment. Those show that the passive approach of fast scanning from side to side is more effective when uncertainty is high. This is reasonable since high

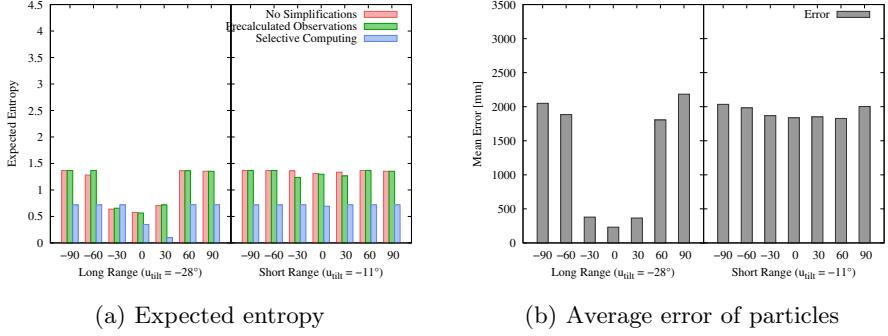


Fig. 2. The expected entropy and the real error of particles after executing an action on a penalty kick position facing the center circle

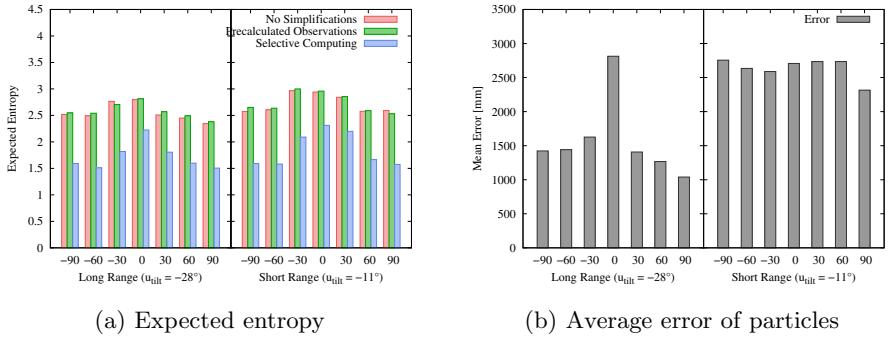


Fig. 3. The expected entropy and the real error of particles after executing an action on the removal-penalty's return position, i.e. on the side line facing the center circle

uncertainty deprives the reasoning attempt of its information to make useful decisions while scanning around for several different observations resolves the uniform distribution fastest. After different distinct position hypotheses can be deduced from the belief state the active approach tends to concentrate on the features providing most information optimally reducing the remaining uncertainty to result in a more precise position estimate.

4.2 Dynamic Situations

A final experiment places the robot on the sideline of the field with uncertain prior knowledge of the starting position to be on one of the side lines and the task to walk to a series of target positions (see fig. 4 and 5). Prior localization knowledge is provided since random particle distributions result in random active vision decisions as shown before. The benefit of the active approach is clearly visible especially when only few useful features are in front of the robot. In such

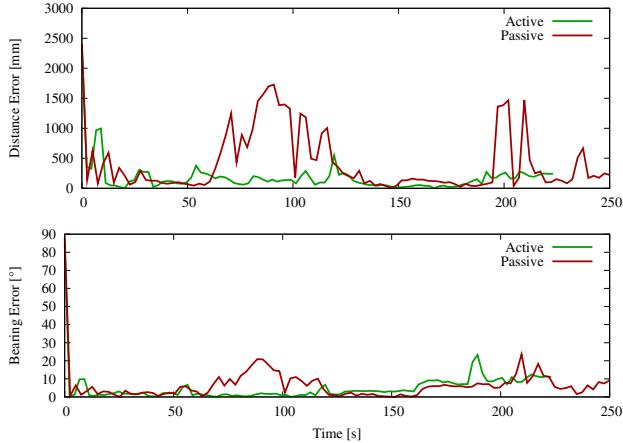


Fig. 4. Position errors (in average 172 mm and 418 mm) and orientation errors (in average 4.97° and 6.14°) of localization while walking to target positions on the field for the active and passive approach, respectively

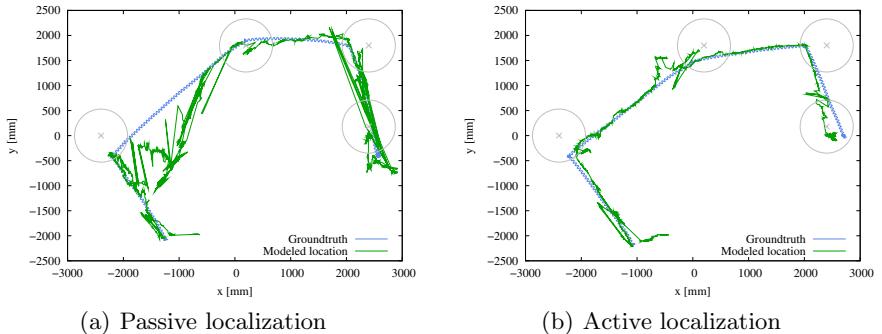


Fig. 5. Localization (green) versus true position (blue) of the robot while walking to a series of target positions

situations most of the scanning motion of the passive approach is wasted on featureless areas.

5 Conclusion

In summary active vision provides the possibility to improve localization in most common situations where the localization algorithm itself already provides some result. For situations of random particle distributions which can be detected with the same entropy criterion a fast scanning motion should be used instead to resolve the high uncertainty by many different observations instead of a few specific ones. Most importantly active vision is also possible and useful even

for platforms with low processing power like the Nao and in environments where relevant features are distributed uniformly as is the case for a Standard Platform League soccer field.

In future work this approach can be extended to also consider dynamic elements of interest like the ball or other robots. Since those are commonly not estimated together with the localization in a single particle filter a multi-modal optimization has to be applied to the decision problem instead of minimization of the entropy of a single particle set.

References

1. Thrun, S., Burgard, W., Fox, D.: Probabilistic Robotics (Intelligent Robotics and Autonomous Agents), 3rd edn. The MIT Press, Cambridge (2005)
2. Burgard, W., Fox, D., Thrun, S.: Active mobile robot localization. In: Proceedings of IJCAI 1997, Morgan Kaufmann, San Francisco (1997)
3. Hernández, M., Cabrera, J., Domínguez, A., Castrillón, M., Guerra, C., Hernández, D., Isern, J.: DESEO: An active vision system for detection, tracking and recognition. In: Christensen, H.I. (ed.) ICVS 1999. LNCS, vol. 1542, pp. 376–391. Springer, Heidelberg (1999)
4. Davison, A.J., Mayol, W.W., Murray, D.W.: Real-time localisation and mapping with wearable active vision. In: ISMAR 2003: Proceedings of the 2nd IEEE/ACM International Symposium on Mixed and Augmented Reality, p. 18. IEEE Computer Society Press, Washington (2003)
5. Porta, J.M., Terwijn, B., Kröse, B.: Efficient entropy-based action selection for appearance-based robot localization. In: Proc. IEEE Int. Conf. on Robotics and Automation, pp. 2842–2847 (2003)
6. Arbel, T., Ferrie, F.: Entropy-based gaze planning. In: Proceedings of the Second IEEE Workshop on Perception for Mobile Agents (1999)
7. Fairfield, N., Wettergreen, D.: Active localization on the ocean floor with multibeam sonar. In: Proceedings of MTS/IEEE OCEANS (2008)
8. Fox, D., Burgard, W., Thrun, S.: Active markov localization for mobile robots. *Robotics and Autonomous Systems* 25, 195–207 (1998)