Autonomous Knowledge Acquisition of a Service Robot by Probabilistic Perception Planning

Robert Eidenberger

Abstract— **This paper presents a concept for active scene modeling by autonomous knowledge acquisition of a service robot in a kitchen scenario. Perceiving complex scenes with multiple objects at arbitrary positions is often difficult from a single measurement due to object ambiguities, object occlusions or environmental influences. The incorporation of several sequential measurements helps to improve scene knowledge.**

In this work a probabilistic active perception framework is developed which plans future sensing actions with respect to uncertainties in the current scene model, unseen spaces and actuation costs of the service robot. Scenes consist of an unknown number of objects, whose poses are modeled in continuous 6D space. The object database contains 100 different household items. The uncertainties in the recognition process and of state transition are probabilistically modeled and basis for planning future perceptions.

The active perception system for autonomous service robots is evaluated in experiments in a kitchen environment. In 200 test runs the efficiency and satisfactory behavior of the proposed methodology is shown in comparison to a random and an incremental action selection strategy.

I. INTRODUCTION

This paper focuses on an active perception system for a household service robot acting in everyday environments. The main scope is the precise scene modeling and sensor action planning to improve scene knowledge. Autonomous knowledge acquisition is essential to enhance the internal representation of the environment of the robot. Figure 1 shows the robot operating in a kitchen scenario. The mobile platform contains two arms with 3-finger hands each for object manipulation. A stereo camera system is used to perceive the environment.

The target operation environments of future service robots are not tailored for automation, rather they are unconstrained, non-cooperative and cluttered. The scenes might be complex containing objects of diverse types in arbitrary 6D-poses, which leads to effects like partial occlusion. High recognition rates, acceptable speed and an accuracy of position $< 1cm$ (as prerequisite for successful manipulation) are required. While many of today's robots typically still live in blocks worlds, it was the goal of this work to address more realistic and hence more complex scenarios. The guiding principles of our approach are a systematic probabilistic handling of the uncertainties of object class and object pose, and the exploitation of the robot's motion capabilities via active perception mechanisms.

R. Eidenberger is with Corporate Research and Technologies, Siemens AG, 81739 Munich, Germany and with the Department of Computational Perception, Johannes Kepler University Linz, 4040 Linz, Austriarobert.eidenberger@gmx.at

Fig. 1. The service robot, operating in an environment with a large number of objects.

The outline of the paper is as follows. The next section presents a brief overview of current approaches to active perception planning. Section III describes the concepts of scene modeling with probabilistic 6D poses. Section IV describes the perception architecture in detail and emphasizes the observation model, the inference model and the planning model. This paper closes with an evaluation of the proposed approach over a large number of experiments in Section V demonstrating satisfactory perception results.

II. RELATED WORK

Literature provides many approaches to active recognition and next best view planning. Surveys on perception planning by Chen[1] and Dutta-Roy[2] list the current state of the art. Here we discuss active perception approaches with respect to their state modeling, the applicability to high-dimensional continuous state spaces their planning concepts and their exploration strategies.

The research in active perception brought up very sophisticated approaches for viewpoint evaluation and next best view planning [3][4][5][6]. They vary in their methods of evaluating sensor positions, in the strategies for action planning and in their field of application. Their main focus is on fast and efficient object recognition of similar and ambiguous objects, but they do not address multi-object scenarios and cluttered environments. These aspects of our work are covered in [7].

The dimensionality of the state space varies among all proposals. Most claim the possible applicability to continuous, high-dimensional domains, but only some successfully prove their theories [5] [8] [9]. The usage in high-dimensional state spaces requires adequate probabilistic pose representations, especially of the orientation. However, nobody tackles the problem of reasonably modeling uncertainties of the orientation for pose determination, which is essential for $6D$ perception.

Active explorative strategies are mostly implemented when dealing with active object reconstruction [10][11][12]. However, they aim at best covering all space but only consider single object scenarios and do not treat the orientation of the objects explicitly.

In this work scenarios are considered, which contain arrangements of several objects, belonging to different as well as to alike classes. Their pose uncertainties are represented by $6D$ multivariate probabilistic distributions. The planning bases on state distributions and explorative criteria.

III. PROBABILISTIC SCENE MODELING

The state space consists of n object hypotheses each represented by a tupel $q^i = (C^i, \phi^i)$. C^i describes its discrete class representation and ϕ^i its continuous pose. In this work a Rodrigues vector representation is used to express the orientation of a 6D pose. All I object-instance-tuples build up the joint state $q = (q^1, q^2, ..., q^I)$. The entities are assumed to be mutually independent. Since q^i represents both discrete and continuous dimensions it will be further considered as a mixed state. The dimension of the state space varies as the perception process proceeds with recognizing new object instances.

In order to describe uncertainties of object hypotheses we use the multivariate Gaussian mixture distribution of the form

$$p(p^i|C^i) = \sum_{k=1}^{K} w_k^i \mathcal{N}(p^i|\mu_k^i, \Sigma_k^i), \qquad (1)$$

to describe the pose accuracy as a complex, multi-peaked distribution for a given object class. K is the number of Gaussian kernels. The mixing coefficient w_k^i denotes the weight of the mixture component with $0 \leq w_k^i \leq 1$ and $\sum_{k=1}^{K} w_k^i = 1$. μ_k^i is the mean and Σ_k^i the covariance of kernel k. The probability distribution over the class $P(C^i)$ is discrete and is described by a histogram over the object classes. $p(q^i)$ is defined as the product of $P(C^i)$ and $p(\phi^i|C^i)$. More on the modeling of pose uncertainty to represent 6D object poses can be found in [13]. The application on the Rodrigues representation is explained in [14].

IV. ACTIVE PERCEPTION ARCHITECTURE

Additional observations from different viewpoints can - via a fusion process - reduce uncertainties with respect to class or pose of objects in the scene or with respect to unexplored regions. Our active perception component determines possible best next views, even in the case of partial occlusions. The framework for selecting the best future action policy π is schematically illustrated in Figure 2. It consists of the *Observation model*, the *Inference model* and the *Planning model*.

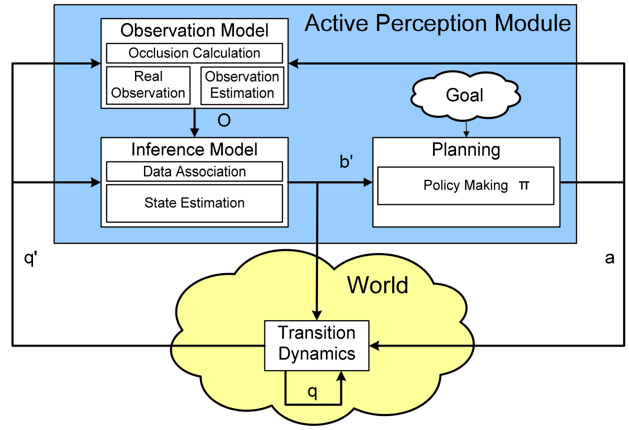

Fig. 2. Active perception framework

The observation model provides the measurement data $O_t(a_t)$ for state estimation. The expected observation is predicted from the chosen sensing action and the predicted state distribution after the transition update. For more accurate observation prediction, object occlusions in the multi-object scenarios are calculated. The measurement likelihood $p(O_t(a_t)|q')$ as a probabilistic representation of the observation is fused in the state estimation process with current scene information.

Perception planning reasons over estimated belief distributions $b_t(q')$ for finding the best action policy in order to reduce the state uncertainty. In this paper we only consider sensor positioning at different viewpoints as sensing actions. The reward function of the planning algorithm takes into account the expected information gain and the estimated costs of the robot motion required. Proposed camera poses are communicated as a *Wish list* to the robot system control for final decision based on geometric-kinematic accessibility.

All these components of the active perception module are explained in detail in the following sections.

A. Inference Model

This work uses the Bayesian state estimator and considers uncertainties in the state transition and in the measurement for state estimation. The probability distribution over the state

$$b_{t-1}(q) = p(q|O_{t-1}(a_{t-1}), ..., O_0(a_0)) \qquad (2)$$

is the a priori belief given previous sensor measurements $O_{t-1}(a_{t-1}), ..., O_0(a_0)$.

The posterior distribution $b_t(q')$ is calculated according to Bayes' rule by updating the prior after the state transition, which is described by its probability $p_{a_t}(q'|q)$, with the new observation $O_t(a_t)$

$$b_t(q') = \frac{P(O_t(a_t)|q') \int_q p_{a_t}(q'|q) b_{t-1}(q) dq}{\int_{q'} P(O_t(a_t)|q') \int_q p_{a_t}(q'|q) b_{t-1}(q) dq dq'}. \qquad (3)$$

The rules of probability, the Markov assumption and the theorem of total probability are applied to derive this expression. The observation $O_t(a_t)$ is assumed to be conditionally independent of previous measurements. For details refer to [15].

Data association is accomplished by combining a global nearest neighbor (GNN) approach and geometry-based data association to find corresponding measurement components. Both build up association tables. GNN data association uses the Mahalanobis distance measure to probabilistically compare Gaussian kernels of pose distributions. This is only applicable for components of the same object classes. Geometry-based data association accomplishes the association task over classes by checking object constellations for physical plausibility, meaning object instances must not intersect. If the entries of the association tables are within the validation gate, the corresponding measurements are associated. Otherwise, unassigned measurements establish a new object instance distribution which is fused in the Bayes' update with a uniform prior, resulting in an increase of the dimension of the joint state.

B. Observation Model

The observation model aims at estimating the observation likelihood $P(O_t(a_t)|q')$ for the current measurement $O_t(a_t)$. Under the assumption of using interest point detectors this observation can be expressed as the detection of a set of N features

$$O_t(a_t) = \{f_1(a_t), ..., f_N(a_t)\}, \tag{4}$$

as a subset of all database features. These features are considered to be the currently visible interest points.

We generate this set of features explicitly when predicting an observation, where we simulate the measurement. Feature characteristics and occlusion events are considered [7]. While for a real measurement the set of features is acquired directly from the detector, during the simulation of the observation we estimate the visibility of a features based on the current scene knowledge. Given the set of expected visible features, $P(O_t(a_t)|q')$ is computed by applying the naive Bayes rule and assuming the features to be conditionally independent:

$$P(O_t(a_t)|q') = \prod_j P(f_j(a_t)|q'). \tag{5}$$

C. Perception Planning

Sequential decision-making consists of the processes of evaluating future actions and finding the best action sequence with respect to a specific goal. The probabilistic planning concept is realized in form of a partially observable Markov decision process as proposed in [15]. The probabilistic planner reasons by considering information theoretic quality criteria of the expected belief distribution $b_t^{O_t(a_t)}(q')$, which is abbreviated by b' in the following equations, and control action costs. The objective lies in maximizing the long term reward of all executed actions and the active reduction of uncertainty in the belief distributions. The value function

$$V_t(b') = \max_{a_t}\left(R_{a_t}(b') + \gamma \int V_{t-1}(b') P(O_t(a_t)|q') dO_t\right) \tag{6}$$

with $V_1(b') = max_{a_t} R_{a_t}(b')$ is a recursive formulation to determine the expected future reward for sequencing actions. γ denotes the discount rate for penalizing later actions and $R_{a_t}(b')$ is the reward. The continuous domains and the high-dimensional state spaces make the problem intractable. As the value function is not piecewise linear, it is evaluated at specific positions, which demands the online calculation of the reward for these specific actions and observations.

The control policy

$$\pi(b') = \operatorname*{argmax}_{a_t}\left(R_{a_t}(b') + \gamma \int V_{t-1}(b') P(O_t(a_t)|q') dO_t\right) \tag{7}$$

maps the probability distribution over the states to actions. Assuming a discrete observation space the integral can be replaced by a sum.

The prospective action policy π is determined by maximizing the expected reward

$$R_{a_t}(b') = \sum_j \alpha_j E_{rel}(R^j_{a_t}(\cdot)) \tag{8}$$

with

$$E_{rel}(R^j_{a_t}(\cdot)) = \frac{E\big(R^j_a(\cdot)\big) - \min_{a_t} E\big(R^j_a(\cdot)\big)}{\min_{a_t} E\big(R^j_a(\cdot)\big) - \max_{a_t} E\big(R^j_a(\cdot)\big)} \tag{9}$$

which relates the relative expected values of several quality criteria j with the respective relation factor α_j.

In this work three different criteria are incorporated, one based on the belief distributions of hypotheses ($R^B_{a_t}(b')$), another on the costs for executing an action ($R^C_{a_t}(b')$) and one on the exploration state of the environment ($R^E_{a_t}(b')$). All three are explained in the following.

1) Information gain from belief distributions: In perception problems the quality of information is usually closely related to the probability density distribution over the state space. The information theoretic measure of the differential entropy is suitable for determining the uncertainty of the belief distribution. Thus, the expected reward from reducing the uncertainty of the state estimates is determined from the expected information gain expressed by the difference of the differential entropies of prior and posterior distribution:

$$E(R^B_{a_t}(q)) = h_{b_t}(q'|O_t(a_t)) - h_{b_{t-1}}(q), \tag{10}$$

Since the computation of the differential entropy both, numerically or by sampling from parametric probability density distributions is costly in terms of processing time, the sum of the upper bound estimates over the object instances

$$\begin{aligned} h^U_{b_t}(q'|O_t(a_t)) &\geq h_{b_t}(q'|O_t(a_t)) \\ &= \sum_i \sum_{k=1}^{K^i} w^i_k[-log w^i_k + \frac{1}{2} log((2\pi \exp)^D |\Sigma^i_k|)] \end{aligned} \tag{11}$$

is used to approximate and determine the expected benefit [16]. D denotes the dimension of the state, $|\Sigma_k|$ denominates the determinant of the kth component's covariance matrix.

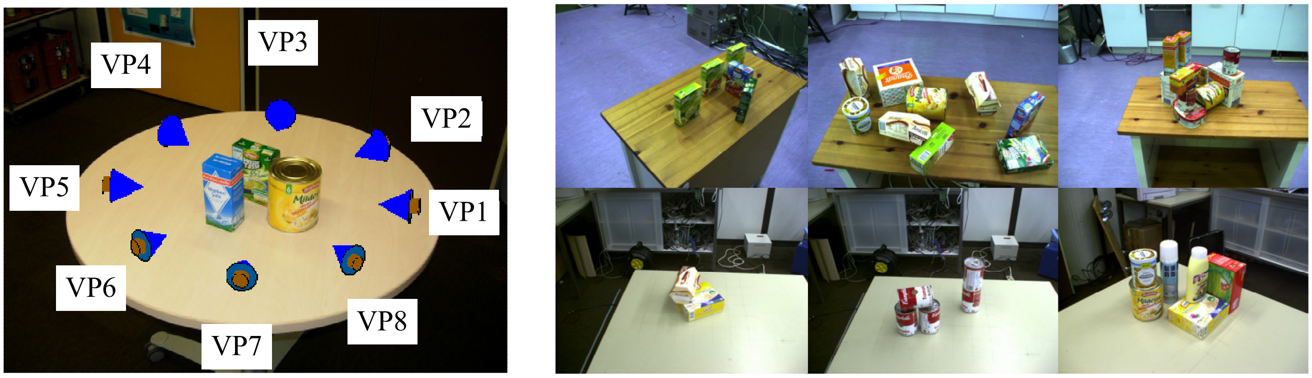

(a) Sample viewpoint arrangement

(b) Sample object constellations

Fig. 3. Sample scenes for experimental evaluation. The objects are selected from the large object database and arbitrarily positioned.

2) Active exploration: Active exploration aims at maximizing the expectation of the reward $R^E_{a_t}(x)$. The state x describes a single volumetric element of the total volume X in the scene. In order to model this volume a grid-based approach in form of a probabilistic octree occupancy grid is used to represent the 3D environment of the translational domain [17]. Each cell in this grid - or also denoted as voxel - has a probabilistic value assigned for its possible stage. The belief $b_t(x)$ is the measure for the degree of exploration. While for the probability $b_t(x) = 1$ the volumetric element x is considered as occupied, for $b_t(x) = 0$ it is entirely free. $b_t(x) = 0.5$ signifies that the state is unknown or unexplored. Initially all probabilities are set to the state unexplored.

In analogy to the determination of the information gain for the hypotheses distributions, the information theoretic measure of the discrete entropy is suitable for determining the uncertainty associated with x[18]. The expected reward from exploration $E\big(R^E_{a_t}(x)\big)$ is set equal to the difference of the prior and posterior entropy of the unexplored volume

$$E\big(R^E_{a_t}(x)\big) = H_{b_t}(x) - H_{b_{t-1}}(x). \qquad (12)$$

3) Costs for executing actions: The execution of each action is associated with the assessed costs $R^C_{a_t}(b')$ from the movement costs of the robot. Thus, the sensor displacement costs are calculated, similarly to [19], by taking into account the robot joints and the drive trajectory.

V. EXPERIMENTS

In this section the proposed approach is compared with two other strategies, random viewpoint selection and an incremental strategy, where the robot either moves clockwise or counterclockwise around the table. The characteristics of the random strategy are plotted in grey in all sequencing figures, of the incremental strategy in orange. For the proposed approaches three different strategies are selected, one for n-step planning (red) and two using a greedy planning horizon, where one has explorative behavior (green). The curves of the simple 1-step strategy are drawn in blue. The evaluation bases on 200 experiments with perceptions of different scenes with varied complexity of up to 10 objects. Figure 3 shows a sample viewpoint constellation of an arrangement of 8 circularly aligned sensing actions and some sample scenes of the experimental series.

Before demonstrating the perception results, the recognition principle which bases on SIFT interest points is described.

A. Object Recognition Principle

Object recognition comprises the tasks of object class identification and pose determination. The challenges of dealing with realistic scenes imposes high demands on the precision of 6D object localization. In this work Lowe's SIFT algorithm [20] is used for determining local, scale-invariant features in images. The application of the SIFT algorithm on stereo images from a camera pair enables the calculation of 6D object poses. This appearance- and model-based approach consists of two separate stages, model generation and Object recognition and pose determination.

Model generation is an off-line process, where the object database is established by determining essential information from training data. We consider a set of 100 household items of different or alike appearance. The challenges lie in the reasonable acquisition and efficient processing and storing of large data sets.

Object recognition and pose determination aims at satisfactory object classification results, low misclassification rates and fast processing. The precise detection of the object pose allows the accurate positional representation of objects in a common reference frame which is essential for proposed planning approach.

Both steps are extensively explained in [21]. The measurement uncertainty is determined from a series of recognitions. The shape of the covariances is assumed to be identical for all objects, their size is dependent on the number of detected interest points and the viewing direction.

B. Number of Iterations and Average Costs

This analysis examines the strategies with respect to the number of iteration steps and actuation costs. The results are plotted in Figure 4a), which compares the average number

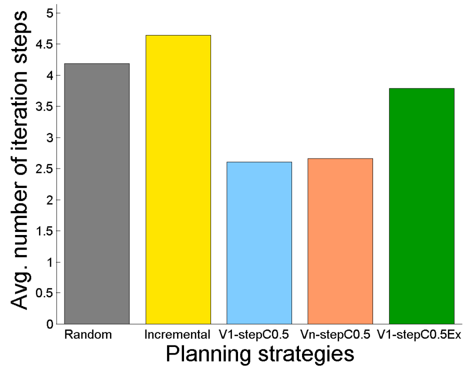

(a) Average number of observation-planning iterations performed for each experiment.

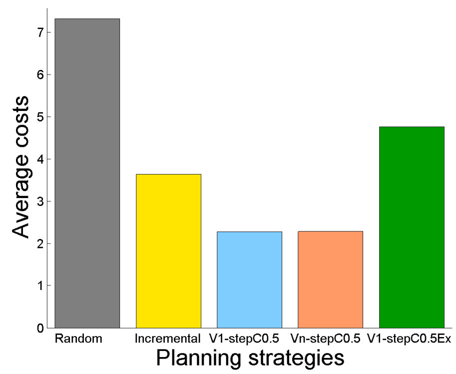

(b) Average robot movement costs when executing a viewpoint change.

Fig. 4. Comparison of the proposed approach with the random and incremental strategy with respect to their number of action-planning combinations and their execution costs.

of required iteration steps to complete the task for each strategy. Despite that the strategy with exploration is weakest among the proposed ones, it still outperforms the other approaches. Less iteration steps signify that less perception-planning tasks have to be performed and less robot actuation is necessary.

Another crucial aspect are the costs, which are involved for each planning concept. According to Figure 4b) which depicts the average costs for each experiment the random strategy with many large action steps is worst. The incremental planning approach is slightly better than the proposed with exploration in the average costs as the step width is very small. The 1-step and n-step planning strategies without explorative behavior are most efficient with respect to the perception and actuation costs.

C. *Object Recognition Rates*

The quality of the perception sequences is evaluated in form of recognition rates. Generally we have to look at different rates, one with respect to the total number of objects in the scene the other regarding all visible object. Generally, the latter rate is higher for one-shot recognition with an average of 72% for detecting the object class under partial occlusions as the number of reference objects is smaller. When referencing the ground truth data of all objects in the scene - even fully occluded and invisible ones - about 60% of these objects are properly recognized with respect to their class properly. When evaluating the class and pose accuracy, which has to be within a certain range, this rate drops to about 38% after the first observation. These rates are fairly constant for all strategies as the initial viewpoint is arbitrarily chosen and the same recognition methodology is applied.

The diagram in Figure 5 shows the recognition rates for the different strategies over the perception-planning iterations. As the approach with exploration explicitly aims at discovering unexplored areas, it is of no surprise that it outperforms all others in the average class detection rates. In contrast, the planning criterion, which bases on belief distributions, targets on improving pose estimates and differentiates objects according to their object class. One can see the slightly steeper slope in the pose recognition rate for the 1-step and n-step planning strategies. Generally, the recognition rates are fairly similar for all approaches. Due to the higher number of observations which are performed by the random and incremental strategy, their recognition rates grow at later iterations, while the proposed approaches often terminate earlier.

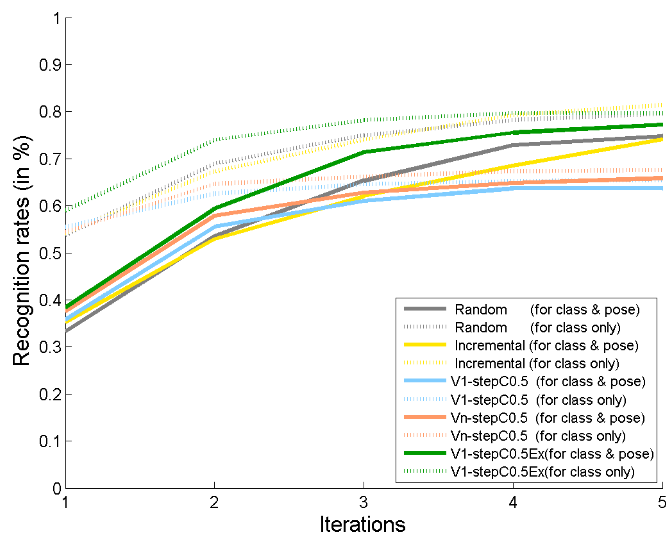

Fig. 5. Class detection rate and recognition rate (where in addition to the correct class the pose deviation from ground truth data is minimal also) with respect to all objects in the scene. Invisible or occluded objects are included in the ground truth basis.

D. *Explorative Behavior*

Explored and unseen volumes are explicitly modeled. The effect of the planning strategies on these volumes are plotted in Figure 6. As the explorative strategy explicitly aims at reducing the unexplored volume it performs best. The strategies with 1-step and n-step planning horizon which do not intentionally reduces the volumetric uncertainty terminate with an average of about 90% of all volume explored, the other strategies explore almost all volume on average. The average explored volume for the incremental and random

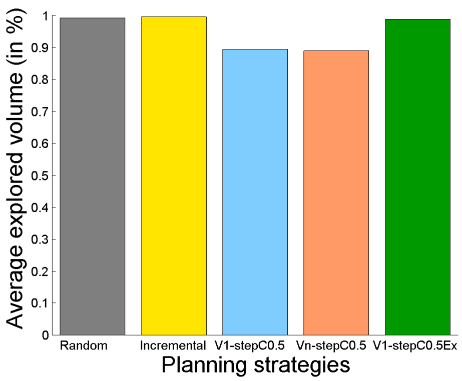

Fig. 6. Variation of the total explored volume for the evaluated strategies

strategy is high due to the large number of sequential observations from different viewpoints, though.

VI. CONCLUSIONS AND FUTURE WORKS

A. Conclusions

In this work we presented a concept for autonomous environmental perception. The recognition of objects bases on SIFT interest points, the active planning system uses measures for the recognition uncertainty, the size of unexplored volume and robot actuation costs to improve scene knowledge. The module is integrated into the DESIRE two-arm mobile platform and evaluated on hundreds of test scenes, which are composed of up to 10 objects out of a 100-object database, which are randomly arranged on a table. The system performed well, so we consider our approach as a contribution towards robots that are able to operate successfully and do useful things in real everyday environments.

B. Future Works

In order to extend the possible range of use, future work will have to address other object types (e. g. poorly textured, partially transparent, shiny) which are not covered by algorithms which are based on texture. Additional recognition and localization algorithms and a deliberate fusion of their results will be required. Again probabilistic methods will have to play a decisive role. They also provide the appropriate platform to utilize a-priori knowledge like physical constraints (e.g. objects do not interpenetrate or levitate). Furthermore, the autonomous learning of new object models has to be tackled to enable the service robot to cope with realistic scenes.

REFERENCES

[1] S. Chen, "On Perception Planning for Active Robot Vision," *IEEE SMC Society eNewsletter*, vol. 13, December 2005.

[2] S. Dutta Roy, S. Chaudhury, and S. Banerjee, "Active recognition through next view planning: a survey," *Pattern Recognition*, vol. 37, no. 3, pp. 429–446, 2004.

[3] J. Denzler and C. M. Brown, "Information theoretic sensor data selection for active object recognition and state estimation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 24, pp. 145–157, 2002.

[4] J. Vogel and N. de Freitas, "Target-directed attention: sequential decision-making for gaze planning," in *International Conference on Robotics and Automation*, 2008.

[5] M. Chli and A. J. Davison, "Active matching," in *Proceedings of the European Conference on Computer Vision*, 2008.

[6] K. Deguchi, "An information theoretic approach for active and effective object recognitions," *SICE Journal of Control, Measurement, and System Integration*, vol. 1, pp. 30–39, 2008.

[7] R. Eidenberger, R. Zoellner, and J. Scharinger, "Probabilistic occlusion estimation in cluttered environments for active perception planning," in *Proceedings of the IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, 2009, pp. 1248–1253.

[8] G. Flandin and F. Chaumette, "Visual data fusion for objects localization by active vision," in *Proceedings of the 7th European Conference on Computer Vision*, 2002.

[9] M. Feixas, M. Sbert, and F. Gonzalez, "A unified information-theoretic framework for viewpoint selection and mesh saliency," 2006.

[10] W. R. Scott, G. Roth, and J.-F. Rivest, "View planning for automated three-dimensional object reconstruction and inspection," *ACM Comput. Surv.*, vol. 35, no. 1, pp. 64–96, 2003.

[11] N. A. Massios and R. B. Fisher, "A best next view selection algorithm incorporating a quality criterion," in *Proceedings of the British Machine Vision Conference*, 1998.

[12] D. Sokolov, D. Plemenos, and K. Tamine, "Viewpoint quality and global scene exploration strategies," in *International Conference on Computer Graphics Theory and Applications*, Setúbal, Portugal, February 2006.

[13] W. Feiten, P. Atwal, R. Eidenberger, and T. Grundmann, *Advances in Robotics Research: Theory, Implementation, Application*. Springer, 2009, ch. 6D Pose Uncertainty in Robotic Perception, pp. 89–98.

[14] R. Eidenberger and J. Scharinger, "Active perception and scene modeling by planning with probabilistic 6d object poses," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2010.

[15] R. Eidenberger, T. Grundmann, and R. Zoellner, "Probabilistic action planning for active scene modeling in continuous high-dimensional domains," in *Proceedings of the IEEE International Conference On Robotics and Automation*, 2009, pp. 2412–2417.

[16] M. F. Huber, T. Bailey, H. Durrant-Whyte, and U. D. Hanebeck, "On entropy approximation for gaussian mixture random vectors," in *Proceedings of the IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, 2008.

[17] P. Payeur, D. Laurendeau, and C. Gosselin, "Range data merging for probabilistic octree modeling of 3-d workspaces," in *International Conference On Robotics and Automation (ICRA'98).*

[18] A. Visser and B. A. Slamet, *European Robotics Symposium 2008*. Springer, 2008, vol. 44, ch. Balancing the Information Gain Against the Movement Cost for Multi-robot Frontier Exploration, pp. 43–52.

[19] É. Marchand and F. Chaumette, "Active vision for complete scene reconstruction and exploration," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 21, no. 1, 1999. [Online]. Available: citeseer.ist.psu.edu/marchand99active.html

[20] D. G. Lowe, "Object recognition from local scale-invariant features," in *International Conference on Computer Vision*, 1999, pp. 1150–1157.

[21] T. Grundmann, R. Eidenberger, M. Schneider, and M. Fiegert, "Robust 6d pose determination in complex environments for one hundred classes," in *Proceedings of the 7th International Conference On Informatics in Control, Automation and Robotics*, 2010, pp. 301–308.