

Probabilistic Rule Set Joint State Update as Approximation to the Full Joint State Estimation Applied to Multi Object Scene Analysis

Thilo Grundmann

Michael Fiebert

Wolfram Burgard

Abstract—One essential capability of service robots lies in the identification and localization of objects in the vicinity of the robot. The extreme computational demands of this high-dimensional state estimation problem require approximations of the joint posterior even for small numbers of objects. A common approach to solve this problem is to marginalize the joint state space and to consider object-related state spaces which are estimated individually under the assumption of statistical independence. In practice, however, this independence assumption is often violated, especially when the objects are located close to each other, which leads to a reduced accuracy of this approximation, compared to the full joint estimation. To address this problem, we propose the new method denoted as Rule Set Joint State Update (RSJSU), which features a better approximation of the joint posterior in the presence of dependencies, and thus leads to better estimation results. We present experimental results in which we simultaneously estimate all six degrees of freedom of multiple objects.

I. INTRODUCTION

State estimation in high dimensional domains is challenging due to the exponential growth of the state space and the non-tractable requirements for computational time and memory. In some cases such as multiple target tracking or multiple object localization, the full state is composed of several sub-states and the joint state is often represented by a set of individual states [1][2]. Under the assumption of statistical independence between all entities, this is equivalent to the full state estimation. In reality, however, total independence between the entities is almost never given.

In this paper we consider the question of how to benefit from statistical dependencies between objects in localization. Utilizing such dependencies can significantly enhance the accuracy of the state estimation process. A typical example is given by methods which exploit knowledge about position and height of a table in the scene to reduce the problem of object localization by three dimensions for objects situated on the table [3][4][5]. Whereas such heuristic approaches are quite effective, they are unable to exploit similar dependencies in more complex scenarios, for example when one object is not placed upright on the table. The goal of this paper is to propose a new, more general method which is able to increase the approximation quality in the presence of dependencies. Our method, denoted as Rule Set Joint State Update (RSJSU) considers physical laws that cause the dependencies in the




Fig. 1. Everyday-life scene for a mobile service robot. Lower left image: object localization result for that scene. Lower right image pair: Corresponding posterior of the five framed objects before (upper figure) and after the application of the RSJSU (lower figure). In this scenario, RSJSU significantly improves the estimation of the objects' vertical position.

object states. It is not restricted with respect to object shapes or poses and hence can be applied in more general and real-life scenarios. The paper is organized as follows: In Section II we motivate our method by discussing the problem domain using an intuitive example. Then we describe our approach in detail in Section III. An implementation of a multi-object localization system and its extension with the new method is presented in Section IV. Finally, we present experimental results in Section V and discuss related work.

II. MOTIVATING EXAMPLE

In this section we motivate our work using a simplified example scenario. We use this example to discuss the approximation errors that occur in the presence of dependencies when independent state spaces are considered in contrast to a full joint state space.

The scenario is discrete and one-dimensional: Two blocks are placed onto a larger one (Fig. 2). For both, there is a sensor that measures their position. This scenario features dependencies since due to the hull of the two bricks it is not possible to place both on the same spot. Their size in combination with the physical laws enforce a minimum

This work has partly been supported by the German Federal Ministry of Education and Research (BMBF) under grant no. 01IME01D, DESIRE.

T. Grundmann and M. Fiebert are with Autonomous Systems, Corporate Technology, Siemens AG, D-80200 Munich, Germany

W. Burgard is with the Department of Computer Science, University of Freiburg, D-79110 Freiburg, Germany




Fig. 2. Scenario with toy bricks' states $x_{blue} = 9$ and $x_{red} = 14$.

distance of 4 between the centers of the two bricks. Knowing this “rule”, it is a priori known which constellations of brick poses are impossible (Fig. 3(a)). However, if both poses are considered independently, this information is lost by marginalization.




Fig. 3. Prior of the joint state of both toy bricks (a). The gray area marks the states that are impossible due to the shape of the two bricks. The state which is depicted in Fig. 2 ($x_J = \{9, 14\}$) is colored. The measurement model $p(z_{blue}|x)$ and $p(z_{red}|x)$ for this state is depicted in (b).

Let us assume that the sensors have a zero mean Gaussian measurement model with different variances (Fig. 3(b)). The posterior can be determined using a Bayes filter update which recursively integrates measurements starting with the prior.




Fig. 4. Full joint posterior (gray area = impossible states) (a), marginals of the joint posterior (b), and independently estimated posteriors (c).

Figure 4 shows, that the results of the joint state Bayes filter and the independent Bayes filter are different for the situation depicted in Fig. 2. The posterior for the red brick is more peaked if we jointly estimate the states compared to the situation in which we estimate them independently.

Comparison of Joint and Independent Posteriors: To quantify the approximation quality of the independent estimation we carried out a simulated experiment in the toy block scenario. For every possible combination of block poses we simulated 250 test runs, each consisting of three measurement update steps. We then measured the KL-divergence between the two posteriors.

As Figure 5 shows, the approximation is very accurate for regions in the joint state where the blocks are far away from each other. However, when the two objects are placed closely the approximation error grows considerably. The goal of this work is to develop a method which is able to enhance the approximation quality in such areas.




Fig. 5. Mean KL-divergence (250 simulation runs per possible true joint state) between joint state and independent state estimation for all combinations of brick-states. The gray area corresponds to the states that are impossible due to collisions.

III. RULE SET JOINT STATE UPDATE

In this section we introduce our new method of “Rule Set Joint State Update” (RSJSU), which increases the approximation quality of independent multi-object state estimation.

A. Independence Assumption in Multi-Object State Estimation

Let us suppose that there are n objects and that the full state x , which consists of n sub-states $x = (x^1, x^2, \dots, x^n)$ is estimated using a recursive Bayes filter according to (1).

$$Bel(x_t) = \eta p(z_t|x_t) \int p(x_t|u_t, x_{t-1}) Bel(x_{t-1}) dx_{t-1} \quad (1)$$

Let us furthermore assume that the prior $Bel(x_0)$, the measurement and the prediction models for the individual objects can be specified under the assumption that they are independent of the other objects. Then we obtain the following equations:

$$Bel(x_t^i) = \eta^i p(z_t^i|x_t^i) \int p(x_t^i|u_t^i, x_{t-1}^i) Bel(x_{t-1}^i) dx_{t-1}^i \quad (2)$$

and

$$Bel(x_t) = \prod_{i=1}^n Bel(x_t^i). \quad (3)$$

In many real life situations the assumption of full independence is not justified. However in such a case the independent estimation can be seen as approximation of the true full state estimation process.

Dependencies in real life scenarios occur when objects are close together or stacked. These dependencies originate from the laws of physics which forbid intersection of rigid objects and require balance of power. We define a rule set r to consist of physical laws and 3D models. In the motivating example given above, the rule set r would comprise both bricks' models and the law that they cannot intersect.

B. Modeling the Rule Set

Statistical dependencies between sub-states have to be represented in the full joint state prior. Formally this is

represented in (7) by $p(x | r)$ as prior which is conditioned on the rule set r . The factors α , α' and α'' act as normalizers.

$$p(x | z_{0..t}, r) = \alpha \cdot p(z_{0..t}, r | x)p(x) \quad (4)$$

$$\stackrel{\text{cond.ind.}}{=} \alpha \cdot p(r | x)p(z_{0..t} | x)p(x) \quad (5)$$

$$= \alpha \cdot p(z_{0..t} | x)p(x | r)p(r) \quad (6)$$

$$\stackrel{\text{cond.ind.}}{=} \alpha' \cdot \prod p(z_i | x)p(x | r) \quad (7)$$

This prior can be imagined as mostly equally distributed with certain areas containing zeros. These areas correspond to the physically impossible states. As the pre-calculation of this joint prior for the full joint state in continuous high dimensional domains seems to be intractable, the solution lies in (8) according to which the rules in the $p(r | x)$ can be applied after the incorporation of measurements.

$$p(x | z_{0..t}, r) = \alpha'' \cdot p(r | x)p(x | z_{0..t}) \quad (8)$$

Hence, $p(r | x)$ must only be evaluated for state regions, that bear probability mass after the incorporation of measurements. Comparable to the measurement model, $p(r | x)$ evaluates how a given state complies to a given rule set. In the case of the physical rule, which enforces that objects cannot intersect, the term would be zero for all x that describe a colliding object constellation.

C. Rule Set Joint State Update

The input parameters of the RSJSU are a set of independently estimated probability distributions $p(x^i | z_{0..t})$ over k sub-states that are suspected to be dependent regarding r .

The first step is to construct the conditional joint posterior under the assumption of independence (9). The constellation of sub-states does not change during RSJSU, so $x^{1..k}$ is abbreviated to $x^{[J]}$.

$$p(x^{[J]} | z_{0..t}) = p(x^1 | z_{0..t})p(x^2 | z_{0..t}) \dots p(x^k | z_{0..t}) \quad (9)$$

To obtain $p(x^{[J]} | z_{0..t}, r)$, $p(x^{[J]} | z_{0..t})$ must be updated with $p(r | x^{[J]})$ according to (8).

Given this we can continue the recursive Bayesian estimation in the joint state. In the case that the application needs an estimate of the sub-state, the joint probability distribution has to be marginalized:

$$p(x^i | z, r) = \int_{m \neq i} p(x^{[J]} | z, r) dx^m. \quad (10)$$

D. Invocation of the RSJSU during Runtime

The gain of the RSJSU depends on when it is applied in an application scenario. Hence, the decision when to apply it is based on a trade-off between the expected gain and the costs.

The expected computational costs for the RSJSU depend on the number of sub-states that are considered and on the granularity of the corresponding models. These factors are known at runtime, but the expected gain is unknown and must be estimated. Since the gain grows with growing dependency, an estimation for the dependency could be used as estimation for the gain.

E. RSJSU Applied to the Motivating Example

In section II we compared independent estimation and full joint state estimation using the KL-divergence (Fig. 5). To evaluate the approximation quality of the RSJSU we performed the same evaluation. In particular, we took the results of an independent estimation after the incorporation of three measurement steps, updated it using RSJSU and compared it to the full state estimation using the KL-divergence (Fig. 6). In this scenario where the true joint prior




Fig. 6. Mean KL-divergence (250 simulation runs) between joint state and RSJSU state estimation for all possible brick-states. The scale is the same as in Fig. 5.

is known and used in the RSJSU, the approximation quality is substantially higher. The maximum mean KL-divergence between RSJSU and JE is about 10^{-16} compared to 0.89 for the independent estimation. However, in reality the true prior will never be known exactly. Hence, in a real world example this extreme precision will not be reached.

IV. IMPLEMENTATION

Our real-world application for the use of RSJSU is the DESIRE mobile service robot (Fig. 1) which features the ability to manipulate every-day objects in a household scenario.

First we will describe the robot's current multiple-classes-multiple-instances object localization system. After that we will explain the implementation of RSJSU and the way it is applied within the current system.

A. Multiple-Classes-Multiple-Instances Object Localization

The role of the object localization is to detect and to localize objects of different classes in 6D. The objects may appear in arbitrary numbers and positions and the "scenes" are assumed to be static. In this context, a class is defined as all object-instances that "look exactly the same". Thus, two cans of Campbell's Tomato Soup 10 oz. are of the same class, whereas another tomato soup does not belong to that class. All classes which shall be manipulated are initially known to the system in terms of a model, although the number of instances of one class that can appear in the scenes is unknown.

As the total number of objects is unknown, the dimension of the overall state which is to be estimated is also unknown. However, as the dimension depends linearly on the number of objects, it is large in realistic scenarios.

To fulfill the requirements while still being tractable, we decided to estimate the states of single objects separately, assuming full mutual probabilistic independence between all objects.

Since this approach requires data association, we first describe the method of data association, which assigns the measurements to the separate targets.

After that we describe the state, measurement model, motion model and the Bayes filter itself which estimates the state of those targets.

B. Data Association

We use a maximum likelihood approach with a built-in threshold, which means that a new measurement is either assigned to the object that accounts for the maximum likelihood $p(z | x)$ to have originated this measurement, or a newly created object, when all existing objects' measurement likelihoods fall below the threshold.

Since the localization system that we are using is quite accurate in relation to the average size of the objects, false data associations do not occur in our experiments. Nevertheless in a different scenario this data association approach could be replaced by a more complex approach like the Monte Carlo Joint Probabilistic Data Association Filter [6].

C. Particle Filter for Recursive State Estimation

For the individual object state estimation we use a recursive Bayes filter, to incorporate all measurements that were assigned by the data association over time. This Bayes filter is implemented using the bootstrap particle filter [7] from the Orocos project [8]. As sensor for the object recognition we use a pair of calibrated high resolution (1388 x 1038) Firewire cameras mounted on a pan-tilt unit.

1) *The State:* The state of a single object consists of a class c and a pose ω with reference to the world frame. This seven-dimensional state is called hypothesis h consisting of one discrete and six continuous dimensions: $h = (c, \omega)$. For the pose we use a Rodriguez representation, which means that the first three entries of the pose vector describe its translational part and the last three entries are formed out of the unity rotation axis (e^1, e^2, e^3) and the factor β that determines the rotation angle (11).

$$\omega = (x^1, x^2, x^3, \beta e^1, \beta e^2, \beta e^3) \quad (11)$$

2) *The Prior:* The prior $p(x_0)$ in our system is assumed to be equally distributed over classes and poses.

3) *The System Model* $p(x_t | x_{t-1})$: As we consider only stationary objects in our scenario, the system model consists of a static state transition.

4) *The Measurement Model* $p(z | x)$: Our localization method is based on the sift feature [9]. To retrieve classification and full 6D localization, we use a combination of stereo vision and 3D sift models which are constructed using the Interactive Modeling Center of the Karlsruhe Institute of Technology [10]. Our approach has no restrictions regarding the object shape and is described in [11].




Fig. 7. Translational errors (cm) of the localization method, compared to ground truth; depicted in camera coordinate system with fitted covariance ellipsoid (70%) for x_c, y_c and z_c (vertical axis).

Probabilistically we describe our localization method with the conditional probability $p(z | x_c)$, x_c being a seven dimensional object hypothesis expressed in the camera coordinate system and z being the measured hypothesis which is delivered by the underlying object localization system. The evaluation of the measurement model $p(z | x_c)$ for specified z and x_c first detects whether the object would be visible in the camera and not too far away. If both conditions are met, the measurement observation likelihood is evaluated using the probability distribution shown in Fig. 7. The necessary normalization is automatically done by the particle filter.

D. Rule Set Joint State Update for Monte-Carlo-Represented Multi-Object Scenarios

In our implementation, we restricted the joint state to consist of two objects' states. However it can be extended to a multi object implementation. The physical fact that scenes have to be free of forces to be stationary, formed the rule in our rule set r (comp. (13)). This means that neither intersecting nor flying object are allowed.

The implementation of the RSJSU (Fig. 8) follows the method description in section III. First, the joint state is constructed by combining two samples from each of the originating two probability distributions into a joint sample. This step is repeated until the desired number of joint particles is reached.

Next, all joint particles are weighted with $p(r | x_{12}^{[n]})$ according to (8) on page 3.

When the joint state representation is no longer required, the joint distribution is marginalized back. Using a Monte Carlo representation, this can be done by "cutting" the joint particles. The resulting sample sets form the distributions of the sub-states.

E. Evaluating $p(r | x_{12}^{[n]})$ Using the Bullet Physics Engine

We use the Bullet Physics Engine to find the maximum penetration depth d_c in a scene $x_{12}^{[n]}$. In the case of hovering objects that do not collide we use the simulation to find a valid constellation and determine the distance d_f to the original object's pose. Both values approximate how far a valid constellation is. Since our sample-based representation is only an approximation when using a limited number of particles, we need to soften the hard freedom of forces constraint. Otherwise no valid constellation would ever appear.

Rule set joint state update

$$\bar{\chi}_{12} = \chi_{12} = \emptyset$$

for $n = 1$ to N do

sample $x_1^{[n]} \sim \text{bel}(x_1)$

sample $x_2^{[n]} \sim \text{bel}(x_2)$

$$x_{12}^{[n]} = \{x_1^{[n]}, x_2^{[n]}\}$$

$$w_{12}^{[n]} = p(r | x_{12}^{[n]})$$

$$\bar{\chi}_{12} = \bar{\chi}_{12} + \langle x_{12}^{[n]}, w_{12}^{[n]} \rangle$$

end for

for $n = 1$ to N do

draw i with probability $\propto w_{12}^{[i]}$

$$\chi_{12} = \{ \chi_{12}, x_{12}^{[i]} \}$$

end for

–Joint state established–

marginalize χ_{12} into χ_1 and χ_2

Fig. 8. The rule set joint state update algorithms for Monte Carlo represented probability distributions.

To account for these approximation errors we use a zero mean Gaussian on the maximum of these two distances and determined the standard deviation σ experimentally.

$$d = \max(d_f(x_{12}^{[n]}), d_c(x_{12}^{[n]})) \quad (12)$$

$$p(r | x_{12}^{[n]}) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{d}{\sigma}\right)^2} \quad (13)$$

The particle representation accomplishes the necessary normalization.

V. EXPERIMENTAL RESULTS

The goal of the new method is to selectively increase the accuracy during independent multi object state estimation in the presence of dependencies between objects.

We tested the method in real-life scenes (Fig. 1), where numerous objects were placed on a table. The gain seemed to be significant by visual inspection, however through the absence of ground truth information for such complex scenes we were not able to quantify the benefit.

Therefore we evaluated the gain of RSJSU by comparison to the classical independent method in a test set that contains 403 different scenes of a different kind.

We designed the test scenes such, that we are able to determine the ground truth poses of the objects as basis for the comparison and strong dependencies between the objects constitute the application of RSJSU. The scenes consist of two objects (Fig. 9), one of them being a table with a calibration pattern which can be localized precisely and the other one being one out of seven different grocery items which can be recognized with the object localization method sketched in section IV. These grocery items feature a calibration sheet on the bottom, to allow for accurate placement onto the calibration pattern, so ground truth poses in relation to the calibration pattern could be recorded.

Only one measurement was executed and incorporated according to the Gaussian measurement model in Section IV-C.4. After that, the mean value of the resulting pose estimate




Fig. 9. One test image (out of 403) with a grocery item and a calibration pattern. The ground truth projected onto the object in the right image.

is compared to ground truth. Accordingly, the mean value of the pose estimate after the application of the RSJSU is also compared to ground truth. It has to be noticed, that in the case of the RSJSU this only makes sense when the updated pose estimate is approximately Gaussian again. This is not generally the case, however in our test scenarios it is. All scenarios were chosen to lie within the approximate working range of our mobile service robot covering different absolute distances from the sensor (0.5m-1.2m) and variable object and camera positions.

A. The Test Results




Fig. 10. Three test situations (depicted in world frame): Particle set (translational part only) from original measurement model: red circles with red covariance ellipsoid at 75 percentile; particle set after RSJSU: blue crosses with blue covariance ellipsoid accordingly.

Fig. 10 shows three exemplary intermediate results from the tests. Due to the directed nature of both of the measurement models and the physical rules, the gain depends on how those directions correlate. In Fig. 10(b) which was taken from a downlooking view point the gain is higher due to the steep view point.

To quantify the advantage of our method, the distribution of the Euclidean error of the translational part compared to ground truth is shown in Fig. 11. The histograms show, that the error of the pose estimation has lowered. With respect to the expected values of both distributions our method yields an improvement of about 15%.

As expected from the physical properties of the table, the highest improvement is achieved in the vertical direction (see Fig. 11 c/d).

VI. RELATED WORK

Dependencies between objects in the state have been used to reduce the state dimension of some of the involved objects

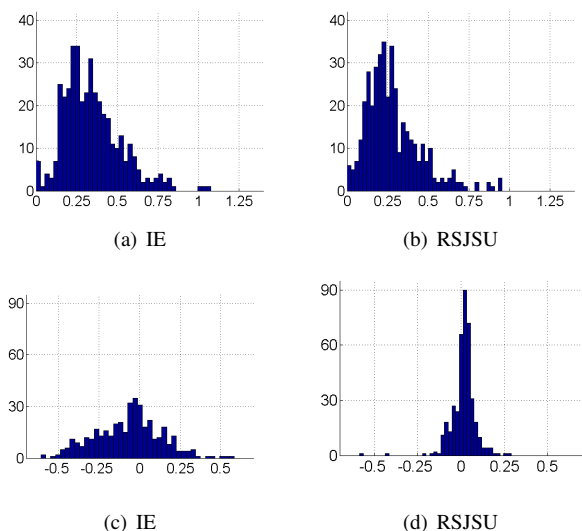


Fig. 11. Upper figures: Histogram of the Euclidean translation error (worldframe) compared to ground truth (cm) (a): mean value 3.43cm (b): mean value 2.82cm . Lower figures: Histogram of the z-axis error (worldframe) compared to ground truth (cm).

in [3], [4] and [5]. The drawback of this approach compared to ours lies in its low generality.

In [12], a map-based motion model $p(x_t | u_t, x_{t-1}, m)$ is presented which is conditioned on the map m to account for violations of (7). That way, the transformation of a particle into a pose that collides with the occupancy map is suppressed. With our approach, exactly this could be done using a joint state motion model. However, as the map-based motion model approach does not explicitly model the path, the same result could be reached with our state dependency approach which would delete the impossible states after the state transition.

Bohlmann [13] presented an outdoor 2.5D localization method which is based on GPS data and enhanced by cadastral data. The sampled posterior is constantly filtered with cadastral data of buildings to delete impossible robot poses. However they assume this data to be exactly known and correct.

Regarding dependencies between objects in the measurement model, Rasmussen and Hager have proposed the Joint Likelihood Filter [2] that explicitly handles occlusion. Since their system estimates a 2D position the depth information of the targets is not included in the state, so it is estimated separately. In the joint state that we propose, the possible occlusion does not have to be sampled since it can be derived directly from the joint state. Kreucher [14] estimates the full joint multi target probability density using a particle filter, and models occlusion directly. Their approach is also based on the insight, that these dependencies occur locally, however they only model dependencies in the measurement model and ignore dependencies in the state which our approach models.

Teather [15] examines the use of physics engines for user interfaces for content creation in virtual reality. They restrict the object motion from 6DOF to physically possible changes,

which f.i. enforces a table to stand on the ground. The aspect of interest here for our paper is that it follows the same basic idea: The physical laws can be used to correct scenes when noisy or lossy input leads to unlikely scenarios.

VII. CONCLUSIONS

In this paper, we proposed the Rule Set Joint State Update as new method to increase the estimation accuracy of independent sub-state estimation in the presence of dependencies in the prior. Compared to the majority of previous methods, our approach features a more general concept to utilize dependencies between objects on the occurrence. Our algorithm has been implemented for a sample-based multi-object localization system. It uses a physics engine to model the physical relations in the prior which cause the statistical dependencies.

Experiments carried out on real data demonstrate the applicability in complex service robotic scenarios. Extensive test performed on a large dataset revealed a substantial gain of precision in the pose estimate obtained with our algorithm in comparison to a standard independent sub-state estimation approach.

REFERENCES

- [1] Y. Bar-Shalom and L. Xiao-Rong, *Estimation and Tracking: Principles, Techniques and Software*. London: Artech House Publishers, 1993.
- [2] C. Rasmussen and G. D. Hager, "Probabilistic data association methods for tracking complex visual objects," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 23, no. 6, pp. 560–576, 2001.
- [3] J. A. Walter and B. Arnrich, "Gabor filters for object localization and robot grasping," in *ICPR*, 2000, pp. 4124–4127.
- [4] J. Zhang, R. Schmidt, and A. Knoll, "Appearance-based visual learning in a neuro-fuzzy model for fine-positioning of manipulators," in *ICRA*, 1999.
- [5] G. v. Wichert, C. Klimowicz, W. Neubauer, T. Woesch, G. Lawitzky, R. Caspari, H.-J. Heger, P. Witschel, U. Handmann, and M. Rinne, "The robotic bar - an integrated demonstration of man-robot interaction in a service scenario," in *Proceedings of the 2002 IEEE Int. Workshop ROMAN, Berlin*, 2002.
- [6] R. Karlsson and F. Gustafsson, "Monte carlo data association for multiple target tracking," 2001.
- [7] A. Doucet, N. de Freitas, and N. Gordon, *Sequential Monte Carlo Methods in Practice*. Springer, New York, 2001.
- [8] K. Gadeyne, "BFL: Bayesian Filtering Library," <http://www.orocos.org/bfl>, 2001.
- [9] D. G. Lowe, "Object recognition from local scale-invariant features," in *International Conference on Computer Vision*, Corfu, Greece, September 1999, pp. 1150–1157.
- [10] R. Becher, P. Steinhaus, R. Zoellner, and R. Dillmann, "Design and implementation of an interactive object modelling system," in *Proceedings of the Joint Conference on Robotics (ISR 2006 ROBOTIK 2006)*, Munich, Germany, 2006.
- [11] T. Grundmann, R. Eidenberger, M. Schneider, and M. Fiebert, "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.
- [12] S. Thrun, W. Burgard, and D. Fox, *Probabilistic Robotics*. MIT Press, Cambridge, MA, 2005.
- [13] K. Bohlmann, P. Biber, and A. Zell, "Using geographical data and sonar to improve gps localization for mobile robots," in *4th European Conference on Mobile Robots (ECMR)*, 09 2009.
- [14] C. Kreucher, K. Kastella, and A. Hero, "Multitarget tracking using the joint multitarget probability density," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 41, pp. 1396–1414, 2005.
- [15] R. Teather and W. Stuerzlinger, "Assessing the effects of orientation and device on (constrained) 3d movement techniques," in *Proc. IEEE Symposium on 3D User Interfaces 3DUI 2008*, 2008, pp. 43–50.