# Towards Dependable Perception: Guaranteed Inference for Global Localization

Anna Petrovskaya, Sebastian Thrun, Daphne Koller, Oussama Khatib
{ anya, thrun, koller, ok } @ cs.stanford.edu
Computer Science Department, Stanford University, CA, USA
**Extended Version, June, 2010**

*Abstract*— Reliable state estimation is an important enabler for robot operation in human environments. Uncertainty and unpredictability of these environments requires global uncertainty problems to be solved for dependable operation. Relative sensors — such as vision, laser and tactile — are common in these applications leading to challenging perceptual problems, for which modern inference methods fail to guarantee accurate estimates. Further, as we show, the reliability of these estimates degrades quickly as initial uncertainty increases. In this paper, our aim is to maximize the amount of information extracted from sensory data, allowing the robot to make the most of its sensors. We present an inference algorithm, which guarantees that all optimal solutions will be found and provides provable error bounds on the approximation of the underlying probability distribution. The approach is based on insight into the sensor model, which is used to guide the refinement process in an adaptive grid algorithm. The approach is applicable to a variety of pose estimation problems with relative sensors. We demonstrate the generality of the approach on the examples of indoor robot localization and tactile manipulation, where it dramatically outperforms state-of-the-art. Empirically, our method increased safety of decision making to 100%. The proposed algorithm also demonstrated logarithmic dependence on desired precision, allowing for efficient high-accuracy estimation. In indoor localization experiments, the approach led to 1mm accuracy of pose estimation based on the commonly used laser range finders. This high accuracy is useful for accurate maneuvering in tight spaces and is sufficient for reliable manipulation of stationary objects of interest within the environment (e.g. door handles, elevator buttons, etc.) It also opens up new potential applications during building construction, inspection and maintenance. In the tactile manipulation setting, the method results in efficient, accurate and reliable 6DOF object pose estimation from tactile data, allowing for reliable manipulation.

## I. INTRODUCTION

Accurate and reliable state estimation is required for safe and dependable robot operation in human environments, where uncertainty is high due to inherent unpredictability. However reliable estimation can be challenging in this setting because the state has to be inferred from non-linear relative sensors. These sensors include the most common types used in robotics: laser, vision, and tactile. In these cases, the underlying probability distribution (the posterior) is highly complex with many discontinuities and narrow peaks. The complexity is due to the properties of the sensors: occlusion boundaries cause discontinuities and high sensor accuracy leads to narrow peaks. Since the sensors are relative, the posterior estimation problem becomes non-trivial whenever

large initial uncertainty must be considered. We will call posteriors exhibiting this type of complexity *high-roughness posteriors* due to their similarity to rough terrain with many sharp transitions and narrow peaks. Although a variety of Monte Carlo and deterministic methods have been developed for non-linear problems [1, 2], the reliability of these methods for high-roughness posteriors degrades as initial uncertainty increases. Methods that can provide guaranteed results for these posteriors are virtually nonexistent.

Due to the challenge posed by high-roughness posteriors, two main paths have been pursued in order to simplify the problem. The first method is to reduce the roughness by smoothing the measurement model. This method is popular in the indoor localization literature [3, 4, 5]. However, sharp transitions in the posterior actually contain very accurate information, which is discarded by smoothing leading to inaccurate estimates and increased ambiguity. The second method is to reduce the initial uncertainty. This method is very popular in visual tracking systems, where the state has to be initialized manually as documented in a recent study of 3D visual tracking literature [6]. However, as the study concludes, methods incapable of dealing with global uncertainty tend to be inherently fragile because they can not recover from tracking failures.

In this paper we present an inference method suitable for moderately dimensional problems, in which the state has to be recovered from global uncertainty based on a set of non-linear relative sensor measurements. The method is a variation of an adaptive grid algorithm [7, 8], in which estimates of posterior variation are used to drive the grid refinement process. Unlike [7, 8], we rely on analysis of the measurement model to produce *sound* bounds on the variation of the posterior. As a result, our approach *guarantees* that all global optima are found and provides *provable* approximation error bounds. The method is capable of handling arbitrarily rough posteriors without discarding information. It can even handle the extreme case of perfect sensors, known to break most popular posterior estimation methods. Thus the method allows the robot to make the most of its sensors by extracting maximum amount of information contained in the sensor data.

We have termed the approach *Guaranteed Recursive Adaptive Bounding* (GRAB). It is applicable to a variety of pose estimation problems with relative sensors. We demonstrate its generality on two basic competency tasks: indoor navigation

and tactile manipulation, where it drastically outperforms state-of-the-art. In our experiments, GRAB increased decision safety to 100%. Empirical evaluation shows that the performance of GRAB scales logarithmically with the desired approximation precision, allowing for efficient high-accuracy estimation. This property is due to the fact that the approach falls within the class of divide-and-conquer methods.

In indoor localization experiments, the approach led to 1mm accuracy of pose estimation based on the commonly used laser range finders. This high accuracy is useful for accurate maneuvering in tight spaces and is sufficient for reliable manipulation of stationary objects of interest within the environment (e.g. door handles, elevator buttons, light switches, etc.) as we have demonstrated in [9]. It also opens up new potential applications during building construction, inspection and maintenance. In the tactile manipulation setting, the method results in efficient and reliable 6DOF object pose estimation with sub-millimeter accuracy, allowing for reliable manipulation.

## II. RELATED WORK

Methods for posterior estimation with non-linear measurements can be divided into deterministic and Monte Carlo classes. Deterministic methods include uniform grid (UG) and, for problems with dynamics, histogram filter [10]. Adaptive grid methods have been developed to improve efficiency by concentrating samples near narrow modes. One of the earliest adaptive grid methods is subregion adaptive integration (SAI), which has been shown by Genz and Kass to outperform Monte Carlo methods in moderate dimensions [7]. SAI selectively subdivides regions based on their estimated contribution to the cumulative integration error. The key difficulty here (and in other adaptive grid methods) is estimating the variation of the posterior within a subregion. SAI estimates the variation by evaluating the function at selected points. This method is error prone and tends to underestimate the variation. Genz and Kass specifically warn that for functions with narrow peaks it is possible to miss the peaks entirely. An adaptive grid algorithm has also been applied to robot localization with laser range finders [8], but similarly to SAI it did not provide any guarantees. A guaranteed adaptive grid algorithm has been developed by Olson for robot localization based on stereo sensors in the context of planetary exploration [11]. However, it did not estimate the full posterior — only its maximum. This approach restricted estimation to just $x$ and $y$ coordinates, obtaining robot orientation from a compass. Moreover, it utilized a likelihood field measurement model [10], which discards negative information. This model is smooth and induces a smooth posterior, but can lead to frequent mis-localizations in the cluttered indoor environments.

Monte Carlo methods include variants of importance sampling and, for problems with dynamics, particle filter [1]. In the mobile robot navigation domain these algorithms are called Monte Carlo localization (MCL) [10]. One of the most widely used adaptive Monte Carlo methods is Adaptive Monte Carlo Localization (AMCL) [12]. It has been implemented in modern mobile navigation suites [4, 13]. AMCL adapts the number of particles over time by considering the KL-divergence of the resulting approximation. Recent work has improved AMCL to better maintain multi-modality during pose tracking by adding spatial clustering [14]. For reliable global localization, all variants of MCL require a large number of samples (at least initially). For this reason, some approaches suggest injecting samples directly from the measurement model [15, 3], although these methods require the availability of a measurement model from which one can easily and efficiently draw samples. A number of smoothing techniques are often applied: e.g. inflation of the measurement noise and/or subsampling of sensor data [3, 10]. Since smoothing discards information, recent approaches improve global localization by learning more sophisticated smoothed models with Gaussian processes [5].

The progress is much slower in tactile manipulation, due to complexity of experimental setup and relatively poor availability of tactile sensors. Posterior estimation literature for this problem relies primarily on variants of particle filters [16, 17, 18]. Since solving a global uncertainty problem in 6DOF with a particle filter is computationally expensive, some authors restrict the problem to 3DOF thus reducing the initial uncertainty [16, 17]. In our prior work [18] we introduced a Scaling Series particle filter (SSPF), which is similar to the technique we describe in this paper. SSPF relies on graduated smoothing to eventually estimate an un-smoothed posterior. However it is not able to provide guarantees.

Belief propagation (BP) and message passing methods have been developed to take advantage of the factorization structure present in some belief estimation problems. These methods can often be more efficient than direct posterior estimation methods and are capable of handling high dimensional problems. Kozlov and Koller proposed a non-uniform discretization BP method (NUBP), which combines BP with adaptive gridding [19]. Similar to SAI, NUBP estimates function variation over subregions by random sampling. As for SAI, this can lead to entirely inaccurate posterior approximations. Recently, the interest in adaptive gridding has been renewed by Isard *et al.* [20]. Isard's approach is similar to NUBP, but requires the integration for BP message computation to be easily tractable, which makes the method unsuitable for the applications we consider. Another recent promising structured method is non-parametric BP (NBP) [21], combining regularization with BP. It turns out that in the problems we consider, local beliefs have even more complex distributions than the posterior itself. Thus, as we show, BP variants perform poorly.

## III. MATHEMATICAL BACKGROUND

In this section we provide mathematical background and introduce the notions necessary for the proposed algorithm.

### A. Problem Statement and Notation

We consider the class of problems where the state $X$ has to be inferred from a set of sensor measurements $\mathcal{D} = \{Y_k\}$. Our goal is to estimate the probability distribution of the

state given the measurements, $p(X|\mathcal{D})$, known as the posterior distribution, which we will denote by $\psi(X)$. Our goal is to approximate the posterior and find all of its modes.

For the general algorithm, we will assume that the state $X$ is a $d$-dimensional vector in a bounded rectangle $R$ in $\mathbb{R}^d$. The measurements are modeled as $K$ random variables $Y_k$, which are drawn independently from conditional probability distributions $p(Y_k|X)$ with domains in $\mathbb{R}^{d_k}$. The conditional probability distributions (CPDs) encode the measurement model and often depend non-linearly on the state $X$. In many applications, the CPDs are naturally given in the log-linear form via measurement energy potentials $v_k : \mathbb{R}^d \times \mathbb{R}^{d_k} \mapsto \mathbb{R}^+$. Then the CPD for $Y_k$ can be written as

$$p(Y_k|X) = \frac{1}{Z_k} \exp\Big(-v_k(X, Y_k)\Big), \qquad (1)$$

where $\frac{1}{Z_k}$ is the normalizing constant. Let $v(X) := \sum_k v_k(X, Y_k)$ be the total measurement energy.

We will primarily focus on problems with global initial uncertainty as this is often the more challenging case in robotics. In this case the prior $p(X)$ is uniform. With this assumption, using Bayes rule and conditional independence of $Y_k$, the posterior can be shown to be proportionate to $\prod_k p(Y_k|X)$. Using the log-linear form of measurement CPDs, we can express this fact as $\psi(X) \propto \exp(-v(X))$. Since the normalization constant is usually unavailable directly, it is more convenient to work with the unnormalized posterior $\pi(X) := \exp(-v(X))$. Then the posterior can be obtained by normalizing: $\psi(X) = \frac{1}{Z}\pi(X)$, where $Z$ is called the partition function. When the prior is non-uniform, we will write $\phi(X)$ to denote the prior, then the posterior is $\psi(X) = \frac{1}{Z}\phi(X)\pi(X)$.

In the indoor localization problem the robot needs to determine its position on a known map $\mathcal{M}$ from laser range measurements $\mathcal{D}$. The map is commonly represented by an occupancy grid, which can be produced using SLAM techniques [10]. The state $X := (x, y, \theta)$ is the robot's pose comprised of map coordinates $(x, y)$ and orientation angle $\theta$. The measurements $\mathcal{D}$ consist of a single scan from a laser range finder. $Y_k := (\rho_k, \alpha_k)$ denotes a single ray in the scan and consists of range $\rho_k$ and bearing $\alpha_k$ components. To interpret the measurements, we rely on the most widely used independent beam (IB) measurement model for range finders [10]. In this model all rays are considered as independent measurements of range to obstacles in the environment, corrupted by Gaussian noise. The expected range to the closest obstacle along ray $k$ is computed by ray tracing on the map. If $\mu_k(X)$ is the expected range along ray $k$, the measurement potential is given by

$$v_k(X, Y_k) := \frac{1}{2\sigma^2}(\mu_k(X) - \rho_k)^2, \qquad (2)$$

where $\sigma^2$ is the Gaussian noise variance.

In the tactile manipulation problem the robot needs to determine the position $X$ of a known stationary object $\mathcal{O}$ based on a set of tactile measurements $\mathcal{D}$. The object is typically represented as a polygonal mesh. The state $X :=$ $(x, y, z, \alpha, \beta, \gamma)$ is the 6DOF pose of the object — including position $(x, y, z)$ and orientation angles $(\alpha, \beta, \gamma)$ — in the manipulator coordinate frame. The measurements $\mathcal{D}$ are obtained by touching the object with the robot's end effector. Each measurement $Y_k := (Y_k^p, Y_k^n)$ consists of the measured cartesian position of the contact point $Y_k^p$ and the measured surface normal $Y_k^n$. To interpret the tactile measurements we use the measurement model proposed in [18]. In this model the measurements are considered independent of each other with both position and normal components corrupted by Gaussian noise. The expected contact point is obtained by maximizing the probability of the measurement. If $\mu_k(X) :=$ $(\mu_k^p, \mu_k^n)$ is the expected contact point for measurement $k$ consisting of contact coordinates $\mu_k^p$ and surface normal $\mu_k^n$, the measurement potential is given by

$$v_k(X, Y_k) := \frac{1}{2\sigma_p^2}||\mu_k^p(X) - Y_k^p||^2 + \frac{1}{2\sigma_n^2}||\mu_k^n(X) - Y_k^n||^2, \qquad (3)$$

where $\sigma_p^2$ and $\sigma_n^2$ are Gaussian noise variances of position and normal measurement components respectively.

### B. Insight into the Measurement Model

Each measurement model contains a wealth of domain knowledge about the application, the sensor, and the measurement process. In order to construct a more efficient inference algorithm, we make two properties of the measurement model available to the algorithm at runtime. The first property is *a relaxation* of the measurement model, representing a more optimistic interpretation of the measurements. The second property is *a strengthening*, representing a more pessimistic interpretation of the measurements. Thus we define $r_k$ to be a *relaxation* and $s_k$ to be a *strengthening* of the measurement potential $v_k$, if for all $X$ in $R$ we have

$$r_k(X) \le v_k(X, Y_k) \le s_k(X). \qquad (4)$$

We define $r(X) := \sum_k r_k(X)$ and $s(X) := \sum_k s_k(X)$. Then for all $X$ in $R$

$$r(X) \le v(X) \le s(X). \qquad (5)$$

From relaxations and strengthenings we obtain bounds on the unnormalized posterior $\pi$. Let $\pi_s(X) := \exp(-s(X))$ and $\pi_r(X) := \exp(-r(X))$. Using (5), for all $X$ in $R$ we obtain

$$\pi_s(X) \le \pi(X) \le \pi_r(X). \qquad (6)$$

To extend the above equation to the case of non-uniform prior, let $\phi_r$ and $\phi_s$ be relaxation and strengthening of the prior, then we have $\phi_s(X)\pi_s(X) \le \phi(X)\pi(X) \le \phi_r(X)\pi_r(X)$.

Note that relaxations and strengthenings defined here are simply lower and upper bounds on the measurement potentials. However, we feel it is helpful to have an intuitive understanding of these bounds, because they are central to the proposed algorithm. For relaxations and strengthenings to be useful, they need to be easy to evaluate on the rectangular grid regions that arise during adaptive gridding. For this reason we construct relaxations and strengthenings that are piece-wise constant over the grid regions.

In the subsections below we show how to build useful relaxations and strengthenings for tactile manipulation and robot localization. These two applications represent two of the most common measurement model types in robotics. Tactile manipulation utilizes a proximity model, which is commonly used for many sensor types: e.g. for stereo [11] and laser range finders (under the name of likelihood fields) [10]. This model is smooth, almost everywhere differentiable and omits negative information. The indoor localization example utilizes the independent beam model, which takes negative information into account. This is a more complex model that results in an erratic discontinuous posterior, but captures more information from the measurement process.

### C. Relaxations and Strengthenings for Tactile Manipulation

Due to the maximization implicit in the definition of the expected contact point, the measurement potential can be computed by maximizing over all the faces comprising the object. Therefore for an object consisting of faces $\mathcal{F}$, $v_k(X, Y_k) = \max_{f \in \mathcal{F}} v_k^f(X, Y_k)$, where $v_k^f$ is a potential for a single face $f$ defined similarly to equation (3). Thus if we construct a relaxation $r_k^f$ and a strengthening $s_k^f$ for the potential of each face, then $r_k := max_f r_k^f$ and $s_k := max_f s_k^f$ will constitute proper relaxation and strengthening of the potential for the whole object.

All that remains is to construct $r_k^f$ and $s_k^f$. An alternative way of computing the measurement potential for a face $f$ is

$$v_k^f(X, Y_k) = \frac{1}{2\sigma_p^2} d(f, Y_k^p)^2 + \frac{1}{2\sigma_n^2} ||n_f - Y_k^n||^2, \quad (7)$$

where $d(\cdot, \cdot)$ denotes the Euclidean distance and $n_f$ denotes the normal of the face $f$.

Consider a grid cell $G$ of object's poses $X$. Let $G_c$ and $G_a$ be the cartesian and angular components of $G$ respectively. We can split up the potential $v_k^f$ into its position and normal components: $v_k^{f,p}$ and $v_k^{f,n}$.

Since the cartesian component of $G$ has no effect on the normal of the face $f$, the difference in normals can increase or decrease by at most $\mathcal{R}(G_a)$, where $\mathcal{R}(\cdot)$ denotes the radius of a body in $n$-space.[1] Hence a strengthening $s_k^{f,n}$ and a relaxation $r_k^{f,n}$ can be defined as $\frac{1}{2\sigma_n^2}(||n_f - Y_k^n|| \pm \mathcal{R}(G_a))^2$ respectively.[2]

For the position component of the potential we need to take into account the contribution from both the cartesian and the angular components of $G$. The cartesian component can increase or decrease the distance to the face by at most $\mathcal{R}(G_c)$. The angular component can increase or decrease this distance by at most $2\mathcal{R}(\mathcal{O})\sin(\frac{\mathcal{R}(G_a)}{2})$ assuming that the object is rotated around the center point, for which all other points of the object lie within $\mathcal{R}(\mathcal{O})$. Hence, we can define a strengthening

$s_k^{f,p}$ and a relaxation $r_k^{f,p}$ by adding and subtracting these quantities[3]: $\frac{1}{2\sigma_p^2}(d(f, Y_k^n) \pm \mathcal{R}(G_c) \pm 2\mathcal{R}(\mathcal{O})\sin(\frac{\mathcal{R}(G_a)}{2}))^2$.

Now we set $r_k^f := r_k^{f,p} + r_k^{f,n}$ and $s_k^f := s_k^{f,p} + s_k^{f,n}$, which define the required relaxation and strengthening of the potential $v_k^f$.

### D. Relaxations and Strengthenings for Robot Localization

Consider a grid cell $G$ of robot's poses $X = (x, y, \theta)$. The $k$-th measurement potential is proportional to the squared error of measured range $\rho_k$ with respect to expected range $\mu_k(X)$ (see Eqn. 2). For the moment assume that $m_k(G)$ is a lower bound on $\mu_k(X)$ for X in G and $M_k(G)$ is an upper bound (we will explain how to obtain these values shortly). Then we can construct a relaxation $r_k$ of the $k$-th measurement potential by underestimating the squared error as follows. If the range reading $\rho_k$ is within $[m_k(G), M_k(G)]$ interval, we set $r_k(X) := 0$. Otherwise we compute the squared error to the closest of the $m_k$ and $M_k$ values:

$$r_k(X) := \frac{1}{2\sigma^2}\min\{(m_k(G) - \rho_k)^2, (M_k(G) - \rho_k)^2\}. \quad (8)$$

A strengthening of the $k$-th potential can be constructed by overestimating the squared error:

$$s_k(X) := \frac{1}{2\sigma^2}\max\{(m_k(G) - \rho_k)^2, (M_k(G) - \rho_k)^2\}. \quad (9)$$

By construction $r_k$ and $s_k$ are constant in $G$ and, as one can easily check, satisfy inequality (4).

All that remains to see is how to obtain $m_k(G)$ and $M_k(G)$. Without loss of generality, let us assume that the range finder is mounted at the center of the robot facing directly forward. A grid cell $G$ of robot's poses creates a cone of possible $k$-th rays (see Fig. 1). We set $m_k(G)$ to the distance between the closest obstacle within this cone and the boundary of $G$. An upper bound $M_k(G)$ on the possible range reading can be computed by walking the piece-wise linear map boundary. If a boundary fragment starts on one side of the cone and finishes on the other, then it completely blocks all rays in the cone from passing through. Hence $M_k(G)$ is attained on this boundary fragment.

For small grid cells, both $m_k$ and $M_k$ can be computed efficiently. For larger grid cells, these values can be propagated from the smaller grid cells comprising the larger ones. All of $m_k$ and $M_k$ values can be pre-computed before the robot needs to localize, because these values are independent of the actual sensor data.

## IV. POSTERIOR APPROXIMATION ALGORITHM AND ANALYSIS

We start with an intuitive description of the proposed algorithm. The algorithm begins by partitioning the state space into large grid cells. Then it iteratively refines grid cells until the desired final resolution is reached. At each iteration, low probability grid cells are pruned to focus computational

---

[1] The radius of a body in $n$-space is the shortest distance such that there exists a point, for which all other points of the body are located within that distance.

[2] The relaxation should be set to zero wherever $||n_f - Y_k^n|| - \mathcal{R}(G_a)$ becomes negative.

[3] Again, if the subtraction of these quantities makes the distance negative, the relaxation should be set to zero.
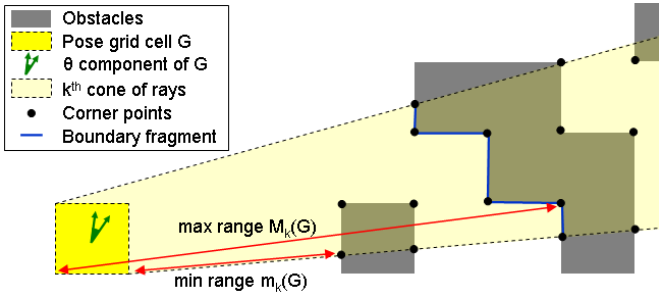
Fig. 1. Computations of min and max range for a cone of ray poses corresponding to a grid cell $G$ of robot poses. Note that the intersections of the occupied space boundary with the cone boundaries are included as corner points.

resources on the high likelihood areas of the state space. The pruning step relies on relaxations and strengthenings. After the final iteration the result is a piece-wise constant approximation of the posterior: zero in all pruned grid cells and equal to each cell's midpoint value in each of the remaining grid cells. We give full algorithm listing in Algorithm 1.[4]

### A. Algorithm Detail

*1) Notation:* Let $\tau^*$ be the desired final resolution — a user specified parameter, which directly controls the degree of refinement. Then the total number of refinement iterations required is $T = log\frac{\delta_1}{\tau^*}$, where $\delta_1$ is the size of region $R$ along the first axis. A finer resolution $\tau^*$ will increase overall running time and improve accuracy of the resulting approximation. Let $vol_*$ be the volume of a grid cell at the final resolution $\tau^*$.

For each iteration $t$, the current set of grid cells will be denoted by $\mathcal{G}^{(t)} := \{G_i | i \in I^{(t)}\}$. All grid cells in $\mathcal{G}^{(t)}$ have the same resolution $\tau^{(t)}$. During pruning step we will prune some grid cells $\mathcal{G}_{prune}^{(t)} \subset \mathcal{G}^{(t)}$ and keep the rest $\mathcal{G}_{keep}^{(t)}$. We will denote the center of a grid cell $G_i$ by $X_i$.

The pruning of low-likelihood grid cells is controlled by a mode sensitivity parameter $\lambda$. It is a user specified parameter, which is the smallest fraction of the maximum value of $\pi$ that is still considered significant. In other words, if $\lambda$ is set to 1%, then all regions where $\pi$ is at least 1% of its maximum value are considered significant.

*2) Initialization:* The initialization of $\hat{\pi}_{max}$ in Eqn. 11 has an effect on the efficiency of the algorithm. The better the estimate the more can be pruned from the very beginning in Eqn. 12. The simplest implementations can set $\hat{\pi}_{max} := 0$ or $\hat{\pi}_{max} := \pi(X)$ for some $X \in R$. A more sophisticated strategy is to run a greedy version of Alg. 1, where $U_{max}^{(t)} := \max_i U_i^{(t)}$ can be used instead of $\hat{\pi}_{max}$ in Eq. 12. Then set $\hat{\pi}_{max} := \max_i \hat{\pi}(X_i)$ based on the posterior estimate produced by the greedy run. The greedy run can not guarantee that a mode will be found, but it tends to provide a good initial estimate for $\pi_{max}$ very quickly.

*3) Bounding Step:* Note that we adaptively change the relaxations and strengthenings from one refinement to the next. In practice $r^{(t)}$ and $s^{(t)}$ are chosen so that the bounds on $\pi_r^{(t)}$

---

[4]In this section we focus on the case of uniform prior. For the non-uniform prior case $\pi_r(X)$, $\pi(X)$, and $\pi_s(X)$ need to be replaced by $\phi_r(X)\pi_r(X)$, $\phi(X)\pi(X)$, and $\phi_s(X)\pi_s(X)$ respectively throughout the section.

and $\pi_s^{(t)}$ are easy to compute. If $L_i^{(t)}$ is a lower bound on $\pi_s^{(t)}$ in a grid cell $G_i$ and $U_i^{(t)}$ is an upper bound on $\pi_r^{(t)}$ in $G_i$, then due to (6) $L_i^{(t)}$ and $U_i^{(t)}$ are also bounds on the variation of $\pi$ within $G_i$:

$$L_i^{(t)} \leq \pi(X) \leq U_i^{(t)}. \tag{10}$$

*4) Pruning Step:* Inequality (12) ensures that the total probability mass discarded after all $T$ refinements is at most $\lambda$ portion of the estimated mass of the "heaviest" grid cell. In practice, for peaked posteriors the value of $\lambda$ can be set quite low, because $\pi$ is virtually zero everywhere except in the vicinity of the peak.

*5) Error Computations:* We have two sources of approximation error in our algorithm. $\varepsilon_{prune}$ upper-bounds the error due to the probability mass of grid cells we pruned. $\varepsilon_{keep}$ upper-bounds the error due to the variation of $\pi$ in the grid cells kept at the final refinement stage. Thus $\varepsilon$ upper-bounds the total $L_1$ error between $\pi$ and its approximation $\hat{\pi}$.

### B. Approximation Analysis

Two important properties of the algorithm can be proven:
1) GRAB will find *all modes* of the posterior.
2) We can compute a sound bound on the $L_1$ distance between the true posterior and its approximation.

Hence we can be absolutely sure that the algorithm did not miss any global optima and we know how good of an approximation was produced.

First we define what we mean by a mode. Let $\psi_{max}$ be the maximum value of the posterior $\psi$, then we will consider any point $X$ in $R$ to be a *$\lambda$-mode* of the posterior if $\psi(X) \geq \lambda\psi_{max}$ for some user-specified parameter $\lambda$.

**Theorem 1:** *Let $\lambda$ be the mode sensitivity setting of GRAB. Let $X$ be a $\lambda$-mode of the posterior $\psi$. Then $X$ is in one of the grid cells kept at the final iteration of GRAB.*

*Proof:* We argue by contradiction. If $X$ is not in the final set of grid cells, then at some iteration $t$ a grid cell $G_i$ containing $X$ has been pruned. However, as one can easily check, this assumption violates the condition (12). ☐

**Theorem 2:** *Let $\hat{Z}$, $\hat{\pi}$ and $\hat{\psi}$ be the partition function, unnormalized posterior and normalized posterior estimates respectively produced by GRAB. Let $\varepsilon$ be the error bound computed by GRAB at runtime. Then*

$$|Z - \hat{Z}| \leq \|\pi - \hat{\pi}\|_{L_1} \leq \varepsilon \tag{17}$$

*and*

$$\|\psi - \hat{\psi}\|_{L_1} \leq \frac{2\varepsilon}{\hat{Z} - \varepsilon}. \tag{18}$$

The proof follows from the bounds carefully constructed at runtime. We provide details in the Appendix.

## V. EXPERIMENTAL RESULTS

### A. Indoor Robot Localization

Indoor environments contain a variety of obstacles invisible to robot's sensors — glass doors, mirrors, staircases — as

**Algorithm 1** Guaranteed Recursive Adaptive Bounding (GRAB): adaptive grid algorithm for posterior estimation.

**Inputs:** $R$ - state space, $\mathcal{D}$ - data set, $\lambda$ - mode sensitivity parameter, $\tau^*$ - desired final resolution.

$$\hat{\pi}_{max} := init\_pi\_max() \qquad (11)$$

Start with the whole region $\mathcal{G}_{keep}^{(0)} := \{R\}$. Until the desired resolution is reached, repeat:

Refine: Construct $\mathcal{G}^{(t)}$ by splitting each grid cell in $\mathcal{G}_{keep}^{(t-1)}$ in halves along each dimension.

Bound: For each $G_i$ in $\mathcal{G}^{(t)}$ compute upper bound $U_i^{(t)}$ and lower bound $L_i^{(t)}$ on $\pi$ using relaxation $r^{(t)}$ and strengthening $s^{(t)}$ and compute $\pi(X_i)$. Update $\hat{\pi}_{max}$, the maximum value of $\pi(X_i)$ observed by the algorithm thus far.

Prune: Add to $\mathcal{G}_{prune}^{(t)}$ grid cells $G_i$ with the lowest upper bounds $U_i^{(t)}$, as long as:

$$\sum_i U_i^{(t)} vol(G_i) \leq \lambda \, \hat{\pi}_{max} vol_*/T. \qquad (12)$$

Prune the selected grid cells by setting $\mathcal{G}_{keep}^{(t)} := \mathcal{G}^{(t)} - \mathcal{G}_{prune}^{(t)}$.

After the final iteration, estimate the unnormalized posterior:

$$\hat{\pi}(X) := \begin{cases} \pi(X_i), & \text{if } X \in G_i \in \mathcal{G}_{keep}^{(T)} \\ 0, & \text{otherwise.} \end{cases} \qquad (13)$$

Estimate the partition function $\hat{Z} := \int \hat{\pi}$ and the normalized posterior $\hat{\psi}(X) := \frac{1}{\hat{Z}}\hat{\pi}(X)$.
Compute approximation error bounds:

$$\varepsilon_{prune} := \sum_{t=1}^{T} \sum_{G_i \in \mathcal{G}_{prune}^{(t)}} U_i^{(t)} \, vol(G_i). \qquad (14)$$

$$\varepsilon_{keep} := \sum_{G_i \in \mathcal{G}_{keep}^{(T)}} (U_i^{(T)} - L_i^{(T)}) \, vol_*. \qquad (15)$$

$$\varepsilon := \varepsilon_{prune} + \varepsilon_{keep}. \qquad (16)$$

**Outputs:** $\hat{\pi}, \hat{Z}, \hat{\psi}$ - estimated unnormalized posterior, partition function, normalized posterior, $\varepsilon$ - approximation error bound.
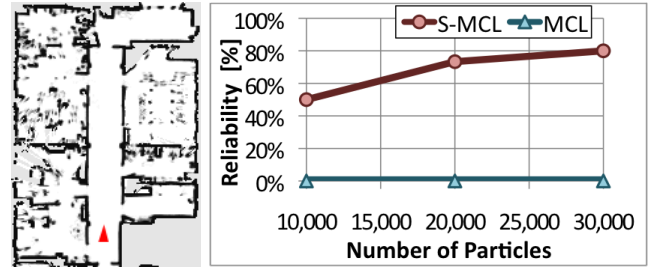


Fig. 2. Smoothing improves reliability of MCL because it decreases roughness of the posterior. Here the reliability of MCL and S-MCL is tested on the 20m x 14m map (left). Red triangle denotes the robot's pose.

each experiment, the localization was performed from a single scan of laser data. The real data was produced by a SICK LMS laser range finder. In the experiments the results are averaged over 100 runs. The error bars in plots represent 95% confidence intervals. For GRAB the different running times are obtained by varying the desired resolution $\tau^*$ with $\lambda = 1\%$.

We compared the performance of the proposed algorithm (GRAB), uniform grid (UG), uniform discretization belief propagation (UBP), non-deterministic belief propagation (NBP, proposed in [21]), subregion adaptive integration (SAI, proposed in [7]), and several variants of Monte Carlo localization. Although for Monte Carlo localization with a single scan of laser data it is more correct to perform a single importance sampling update, in some situations multiple updates can perform better. For completeness we provide both versions: single-update (IS) and 10-update (MCL). For the latter we injected 50cm position noise between updates.

For GRAB, we pre-computed the min-max ranges for grid cells sized 10cm and greater as described in Sect. III-D. The pre-computation only needs to be carried out once for each map as it is independent of actual sensor measurements. For a 70m x 70m map the pre-computation takes 3 minutes. For GRAB, UG, and UBP, we also pre-computed expected range scans for grid cells sized 10cm and greater. This optimization allows grid based methods to be more efficient than Monte Carlo approaches, for which such pre-computations are unsuitable.

The BP algorithms were implemented on a cluster graph satisfying the running intersection property. We assigned $X$ CPD and $Y_1$ CPD factors to the first cluster, and each of the remaining $Y_k$ CPDs to the other $k - 1$ clusters. For UBP, the evidence was applied to the factors prior to propagation. For efficiency reasons, operations during propagation were carried out only on the rows supporting the evidence. Since the evidence was applied prior to propagation, the method is equivalent to integrating measurements one at a time. BP variants performed rather poorly because the belief in this problem is unstructured. While this may seem obvious in retrospect, we still provide the empirical evaluations.

*1) Impact of Smoothing:* In indoor localization the most common smoothing techniques are to increase the measurement noise and to decrease the number of rays considered. To evaluate smoothed MCL (S-MCL) we use 20cm noise and

has been noted in several field studies [22, 23]. Therefore, for safety, it is important for the robot to localize itself on a known map denoting invisible hazards *prior to moving* in the environment. For this reason we evaluate the ability of modern algorithms to reliably estimate the global localization posterior prior to moving.

We performed several sets of experiments with real and simulated robot data in several different environments.[5] In

---

[5]In our experiments we used publicly available maps produced by Cyrill Stachniss (Fig. 2), Mike Montemerlo (Fig. 3), and Ashley Tews (Fig. 4). We thank the map authors for making these maps available to the public.
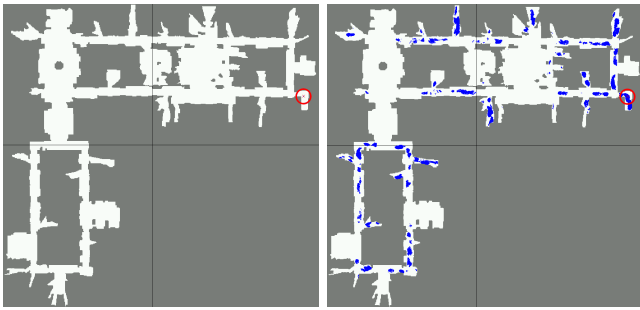
Fig. 3. Smoothing increases ambiguity of localization. Left: without smoothing the resulting posterior has a unique mode (circled). Right: with significant smoothing (50cm noise, 7 rays) the result is highly multi-modal.
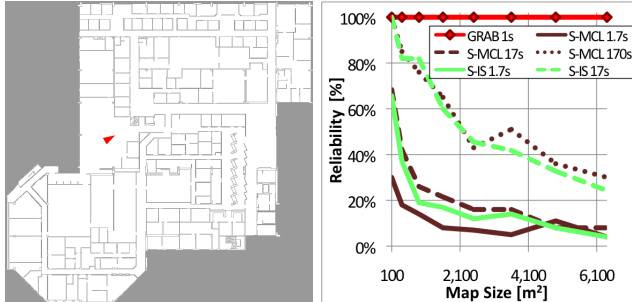


Fig. 4. Dependence of reliability on initial uncertainty (right). Experiments were performed on sub-maps of the 80m x 80m library map (left). Red triangle denotes the robot's pose.
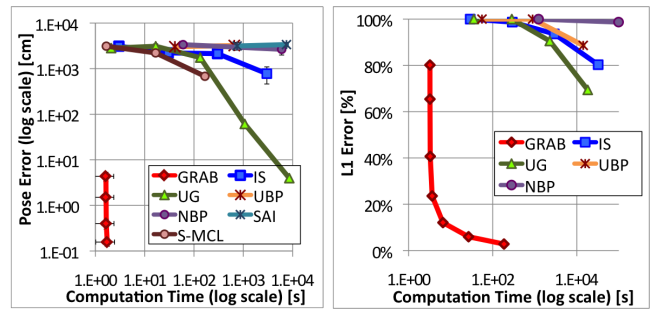


Fig. 5. Accuracy of pose estimation (left) and posterior estimation (right) for global robot localization on the map depicted in Fig. 3. The right plot shows 0.5L1 error, which represents dissimilarity percentage of the approximated posterior vs. ground truth posterior.

61 rays per scan. In this and later experiments, we consider the localization successful if the mean pose is within 1m and 30° of the true pose. Figure 2 shows the reliability of global localization on a 20m x 14m map, where S-MCL exhibits very similar performance to prior art with similar settings [3, 5]. MCL without smoothing (i.e. 1cm noise, 361 rays) fails to solve the problem. Thus smoothing has been vital for solving the global localization problem because MCL is unable to solve the problem without it.

On the other hand smoothing can have a number of negative effects as it discards information contained in the data. Not only can the accuracy go down, but the produced posterior estimate is much more ambiguous than without smoothing (Fig. 3). Under these conditions the robot has to travel to resolve the ambiguity. In addition to being unsafe with respect to invisible hazards, this strategy risks losing the robot's location because MCL does not track multiple modes very well. Moreover, from theoretical perspective smoothed posterior estimates do not converge to the true posterior even as the number of particles tends to infinity.

*2) Dependence on Initial Uncertainty:* Even with smoothing, reliability of MCL drops quickly as map size increases (Fig. 4). These experiments were performed in simulation by taking sub-maps of a larger map. We show results for S-MCL with 1,000, 10,000, and 100,000 particles (with computation time of 1.7s, 17s, and 170s respectively). For comparison, GRAB solves the problem in under 1s (with $\tau^* = 5$cm) and provides guaranteed results.

*3) Localization Accuracy:* To evaluate the accuracy of localization we performed two sets of experiments on the 70m

x 70m map depicted in Fig. 3. In the first set, the robot has to localize after being placed in a random position on the map. These experiments used simulated data with 1cm noise, so that exact reference pose is available for evaluation purposes. The results are shown in Fig. 5 (left), where we plotted the distance of the estimated pose from the reference pose vs. computation time.[6] GRAB demonstrated logarithmic dependence on localization precision, whereas the other algorithms behaved roughly linearly. GRAB was able to recover the reference pose with 1mm accuracy[7] and outperformed other approaches by several orders of magnitude.

The second set of experiments evaluates the accuracy with which the *entire* global localization posterior can be estimated. These experiments were carried out with a single scan of real robot data. The robot's pose is circled in Fig. 3. The reference distribution was obtained using a fine mesh. Prior art approaches were unable to estimate the un-smoothed posterior with any degree of accuracy within reasonable time. The right plot in Fig. 5 shows the results using a smoothed model (25 rays and 10cm noise) for all algorithms and reference distribution. GRAB dramatically outperformed other approaches. The experiments demonstrate that accurate estimation of the posterior takes significantly longer than estimation of only the robot's pose.

*4) Safety During Exploration:* This set of experiments evaluates whether a robot is able to make safe decisions during navigation under conditions of multi-modality. The experiments were carried out in a 70m x 70m simulated museum environment (Fig. 6). The red shaded areas denote safety zones containing glass encased exhibits. The robot was placed randomly in front of the safety zones and had to plan an exploration route using safe maneuvers. Note that this environment has several look-alike areas leading to multi-modal localization posteriors.

First, the robot estimates a global localization posterior. Then it evaluates the proposed action of moving forward by

---

[6]Horizontal error bars for GRAB show 95% confidence in running time for the same setting of $\tau^*$.

[7]Although GRAB was able to recover the pose to within 1mm, one should keep in mind that these experiments were carried out in simulation. To achieve comparable accuracy with real data, accurate environment models are required (see [9] for one such example).
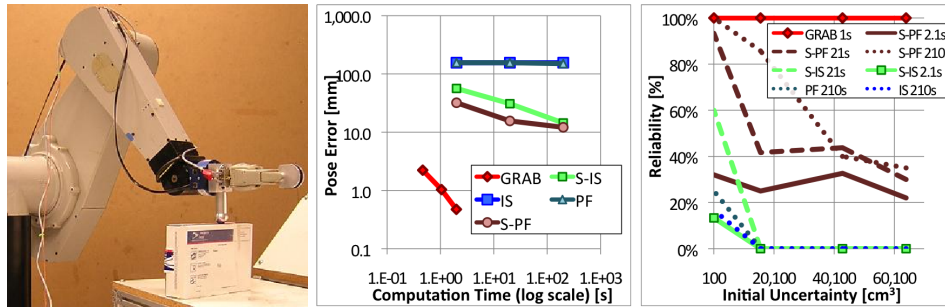
Fig. 7. Tactile sensing for object localization. Left: experimental setup. Center: pose estimation accuracy vs. computation time. Right: reliability vs. initial uncertainty (with unrestricted orientation).
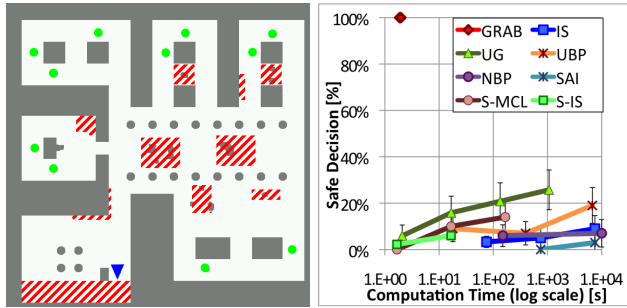


Fig. 6. Decision safety results (right). On the map (left), the blue triangle denotes the robot's pose in one of the experiments. The green circles denote multiple modes of the posterior. Safety zones are shaded in red.

1m based on the estimated posterior. To make a fair comparison of all algorithms, the robot's planner treats all posterior estimates the same. For each experiment, it draws 100,000 samples from the posterior approximation and applies the robot's dynamics model to the samples obtaining a prediction distribution. The prediction distribution is used to compute the expected risk of the maneuver. The risk function is 1 in the safety zones and 0 everywhere else. If the expected risk is higher than 0.01%, the maneuver is considered unsafe. Since the robot is positioned facing safety zones, this maneuver should be identified as unsafe in all experiments.

As expected, GRAB correctly identified all modes of the posterior and thus lead to safe decisions in all experiments. Other algorithms often lead to unsafe decisions.

### B. Tactile Manipulation

In these experiments a stationary object is localized by a robot that explores the object by touching it with its end-effector. The experimental setup consisted of a PUMA manipulator robot equipped with 6D JR3 force/torque sensor at the wrist (see Fig. 7). The sensed object was a rectangular box, for which we constructed a polygonal model using careful ruler measurements. The data sets consisted of five data points taken from different sides of the box.

For un-smoothed versions of the algorithms we used $\sigma_p$ = 1mm and $\sigma_n = 2°$. We compared the proposed algorithm (GRAB), importance sampling (IS), and particle filter (PF). For PF we performed 10 updates with the same data injecting 1cm noise between updates. We also evaluated S-IS and S-PF using smoothing parameters shown to be optimal in [18]: $\sigma_p$ = 1cm and $\sigma_n = 10°$.

*1) Object Localization Accuracy:* First we evaluated how accurately the object could be localized in the robot's workspace. The initial uncertainty in these experiments was 40cm x 40cm x 40cm with unrestricted orientation. The results shown in Fig. 7 (center) are averaged over 100 runs with simulated data, although the algorithms exhibited similar performance with real data. GRAB was able to localize the object to 1mm within about 1s. Other algorithms were unable to get average accuracy better than 1cm even after 3 minutes. It should be noted that sub-centimeter accuracy is often required for successful manipulation [9]. GRAB quickly achieves sub-millimeter accuracy, which we have previously shown to be sufficient for reliable manipulation of objects [18].

*2) Reliability vs. Initial Uncertainty:* To evaluate how reliability depends on initial uncertainty, we varied the uncertainty from 5cm cube to 40cm cube. Unrestricted orientation was used in all experiments. The localization was considered successful if the pose was recovered within 5mm and 5°. The results of 100 runs with simulated data are shown in Fig. 7 (right). GRAB solves the problem in 1s (with $\tau^* = 2$mm) with guaranteed results. We also show results for S-PF with 10,000 to 1,000,000 particles (2.1s to 210s running time) and S-IS with 100,000 to 1,000,000 particles (2.1s to 21s running time). For S-IS and S-PF the reliability degrades quickly with uncertainty. IS and PF performed even worse.

## VI. DISCUSSION AND CONCLUSIONS

We have presented a posterior estimation algorithm that relies on a simple adaptive gridding technique. Yet, in contrast to state-of-the-art, it is able to solve high-roughness posteriors with guaranteed results. The reason for the superior performance is that our algorithm utilizes additional domain knowledge, whereas the other algorithms treat the posterior as a "black box" function. The downside is that GRAB is less general than "black box" approaches and requires additional work to extract the domain knowledge. However, in situations where "black box" algorithms are insufficient, the extra work is very worthwhile. The domain knowledge is captured in a general way in the form of measurement model bounds, which we have given intuitive names of relaxations and strengthenings. Thus our algorithm is applicable to problems where such bounds can be constructed. We have demonstrated the algorithm on indoor navigation and tactile manipulation. We expect that it will extend to other robotic applications with

accurate relative sensors (laser, vision and tactile), although clearly application-specific relaxations and strengthenings will need to be built in each case. Since the class of domain knowledge algorithms is in principle stronger than the class of "black box" algorithms, we hope the demonstrated success will encourage development of algorithms utilizing domain knowledge in a general way.

A number of extensions of the algorithm can be made. For high dimensional problems, the approach can be combined with belief propagation [19, 20] or Rao-Blackwellization [24, 25] — a very promising direction for future work. The refinement strategy could also be tuned to obtain different properties. For example SAI-like refinement strategy could be beneficial when the posterior has regions of similar probability or whenever a better representation of low probability regions is desired. In this paper we have only shown that non-uniform priors could be used in principle, but have not performed any empirical evaluation of such scenarios. The non-uniform priors could carry information about prior measurements and dynamic updates. There are several possible methods for using GRAB to solve dynamic problems. One method is an adaptive histogram filter [8]. Another method is to use GRAB to obtain a representation of the measurement probability distribution for algorithms that sample directly from this distribution [15, 3]. Thus evaluation of how GRAB can be used with non-uniform priors in dynamic applications represents an interesting direction for future work.

Our focus has been solely on posterior estimation and we have left many application aspects outside the scope. For indoor localization, combining the presented algorithm with existing methods for handling high traffic areas, doors and other dynamic environment changes, could provide the community with one of the most reliable and accurate global localization solutions to date.[8] We have shown that GRAB is able to achieve 1mm localization accuracy very quickly even in large buildings. In prior art, it has been shown that laser scans contain sufficient information to achieve such accuracy [27]. In our own prior work, we have shown that 3mm pose estimation accuracy has allowed the robot to reliably manipulate stationary objects not directly observed by its sensors (e.g. door handles, elevator buttons, switches) [9]. Moreover, high localization accuracy is useful for maneuvering in tight spaces (e.g. when passing through a doorway). Of course, accuracy of localization depends on accuracy of the environment model. To achieve high accuracy in [27] laser scans were matched up directly to each other. In [9] a highly accurate polygonal model of a door was used in conjunction with a less accurate grid map. Thus further work on high-accuracy modeling of the environment and sensors will lead to high-accuracy localization, which will open up a host of new applications. In the past, 1mm accuracy has been reserved for stationary manipulators. With high-accuracy localization, mobile robots will be able to aid in building construction,

inspection, and maintenance. They will be able to precisely place fixtures in the environment, accurately drill and paint walls, provide accurate distance measurements, etc. Moreover, high-accuracy localization can easily provide an accurate coordinate transformation between a robot and an off-board sensor (e.g. ceiling mounted camera or another robot's sensor). Hence robots will potentially be able to manipulate even dynamic objects invisible to their own sensors.

In summary, we have presented an adaptive grid method suitable for estimating high-roughness posteriors that result from non-linear relative sensors. The method guarantees that all modes will be found and provides provable approximation error bounds. The strength of the approach comes from utilizing domain knowledge about the measurement process to construct bounds in a general way. We demonstrated the approach on the examples of indoor navigation and tactile manipulation, where it performed significantly better than state-of-the-art and opens up new potential applications.

REFERENCES

[1] A. Doucet and N. De Freitas. *Sequential Monte Carlo Methods in Practice*. Springer, 2001.
[2] M. Evans and T. Swartz. *Approximating Integrals via Monte Carlo and Deterministic Methods*. Oxford University Press, 2000.
[3] S. Thrun, D. Fox, W. Burgard, and F. Dellaert. Robust monte carlo localization for mobile robots. *Artificial Intelligence*, 128(1-2):99–141, 2001.
[4] B. Gerkey, Vaughan R. T., and A. Howard. The player/stage project: Tools for multi-robot and distributed sensor systems. In *ICAR*, 2003.
[5] C. Plagemann. *Gaussian Processes for Flexible Robot Learning*. PhD thesis, University of Freiburg, December 2008.
[6] V. Lepetit and P. Fua. *Monocular model-based 3D tracking of rigid objects*. Now Publishers Inc, 2005.
[7] A. Genz and R. Kass. Subregion adaptive integration of functions having a dominant peak. *Journal of Computational and Graphical Statistics*, 1997.
[8] W. Burgard, A. Derr, D. Fox, and A. B. Cremers. Integrating global position estimation and position tracking for mobile robots: The dynamic markov localization approach. In *IROS*, 1998.
[9] A. Petrovskaya and A.Y. Ng. Probabilistic mobile manipulation in dynamic environments, with application to opening doors. In *IJCAI, Hyderabad*, 2007.
[10] S. Thrun, W. Burgard, and D. Fox. *Probabilistic Robotics*. MIT Press, 2005.
[11] C.F. Olson. Probabilistic self-localization for mobile robots. *IEEE Transactions on Robotics and Automation*, 16(1):55–66, 2000.
[12] D. Fox. Adapting the sample size in particle filters through KLD-sampling. *IJRR*, 2003.
[13] Willow Garage. *ROS — Robot Open Source*. http://www.willowgarage.com/pages/software/ros-platform, May 2009.
[14] Z. Liu, Z. Shi, M. Zhao, and W. Xu. Adaptive dynamic clustered particle filtering for mobile robots global localization. *Journal of Intelligent and Robotic Systems*, 53(1):57–85, 2008.
[15] S. Lenser and M. Veloso. Sensor resetting localization for poorly modelled mobile robots. *ICRA*, 2, 2000.
[16] K. Gadeyne and H. Bruyninckx. Markov techniques for object localization with force-controlled robots. In *ICAR*, 2001.
[17] S.R. Chhatpar and M.S. Branicky. Particle filtering for localization in robotic assemblies with position uncertainty. In *ICROS*, 2005.
[18] A. Petrovskaya, O. Khatib, S. Thrun, and A. Y. Ng. Bayesian estimation for autonomous object manipulation based on tactile sensors. In *ICRA*, pages 707–714, 2006.
[19] A.V. Kozlov and D. Koller. Nonuniform dynamic discretization in hybrid networks. In *Proc. UAI*, pages 314–325. Morgan Kaufmann, 1997.
[20] M. Isard, J. MacCormick, and K. Achan. Continuously-adaptive discretization for message-passing algorithms. In *NIPS*, pages 737–744. MIT Press, 2008.

---

[8]Development of GRAB library is currently underway. Once complete the library will be made available at [26].

[21] E.B. Sudderth, A.T. Ihler, W.T. Freeman, and A.S. Willsky. Nonparametric belief propagation. In *CVPR*, volume 1, page 605, 2003.
[22] W. Burgard, A.B. Cremers, D. Fox, D. Hahnel, G. Lakemeyer, D. Schulz, W. Steiner, and S. Thrun. The interactive museum tour-guide robot. *AAAI*, pages 11–18, 1998.
[23] M. Montemerlo, J. Pineau, N. Roy, S. Thrun, and V. Verma. Experiences with a mobile robotic guide for the elderly. *AAAI*, pages 587–592, 2002.
[24] K. Murphy and S. Russell. *Rao-blackwellized particle filtering for dynamic bayesian networks*. Springer, 2001.
[25] M. Montemerlo. *FastSLAM: A Factored Solution to the Simultaneous Localization and Mapping Problem with Unknown Data Association*. PhD thesis, Robotics Institute, Carnegie Mellon University, 2003.
[26] A. Petrovskaya. *Home Page of Anna Petrovskaya*. `http://cs.stanford.edu/people/petrovsk`, May 2010.
[27] A. Censi. On achievable accuracy for pose tracking. In *ICRA*, 2009.

## APPENDIX

*Proof of Theorem 2:* To see the left hand side of inequality (17), consider that $|Z - \hat{Z}| = |\int \pi - \hat{\pi}|$ and $\|\pi - \hat{\pi}\|_{L_1} = \int |\pi - \hat{\pi}|$. Since $|\int f| \leq \int |f|$ for any integrable function $f$, we can obtain the inequality by setting $f := \pi - \hat{\pi}$.

To see the right hand side of inequality (17), let $R_{prune}$ be the union of all grid cells we pruned and $R_{keep}$ be the union of grid cells we kept at the last iteration. By definition $\|\pi - \hat{\pi}\|_{L_1} = \int |\pi - \hat{\pi}|$. We can split up this integral into a sum of two terms: an integral over $R_{prune}$ and an integral over $R_{keep}$. Since $\hat{\pi}$ is zero everywhere in $R_{prune}$, the first term is simply $\int_{R_{prune}} \pi$, which is bounded from above by $\varepsilon_{prune}$.

To bound the integral over $R_{keep}$, consider one of the grid cells $G_i$ kept at the last iteration. Everywhere in $G_i$, both $\pi$ and $\hat{\pi}$ are bounded by the variation bounds $U_i^{(T)}$ and $L_i^{(T)}$. Hence, $|\pi - \hat{\pi}| \leq U_i^{(T)} - L_i^{(T)}$ everywhere in $G_i$. Integrating this inequality over all grid cells kept at the last iteration and taking into account Eqn. (15), we obtain $\int_{R_{keep}} |\pi - \hat{\pi}| \leq \varepsilon_{keep}$. Thus $\int |\pi - \hat{\pi}| \leq \varepsilon_{prune} + \varepsilon_{keep} = \varepsilon$.

Lastly, to obtain inequality (18), we have the following derivation:

$$
\begin{aligned}
|\psi - \hat{\psi}| &= \left| \frac{1}{Z}\pi - \frac{1}{\hat{Z}}\hat{\pi} \right| = \left| (\frac{1}{Z}\pi - \frac{1}{Z}\hat{\pi}) + (\frac{1}{Z}\hat{\pi} - \frac{1}{\hat{Z}}\hat{\pi}) \right| \\
&= \left| \frac{1}{Z}(\pi - \hat{\pi}) + (\frac{1}{Z} - \frac{1}{\hat{Z}})\hat{\pi} \right| \\
&\leq \frac{1}{Z}\left| \pi - \hat{\pi} \right| + \left| \frac{1}{Z} - \frac{1}{\hat{Z}} \right|\hat{\pi} \\
&= \frac{1}{Z}\left| \pi - \hat{\pi} \right| + \frac{|Z - \hat{Z}|}{Z\hat{Z}}\hat{\pi}.
\end{aligned}
$$

Integrating the obtained inequality over $R$, we have

$$
\begin{aligned}
\int |\psi - \hat{\psi}| &\leq \int \frac{1}{Z}\left| \pi - \hat{\pi} \right| + \int \frac{|Z - \hat{Z}|}{Z\hat{Z}}\hat{\pi} \\
&\leq \frac{1}{Z}\varepsilon + \frac{1}{Z}|Z - \hat{Z}| \leq \frac{2\varepsilon}{Z} \leq \frac{2\varepsilon}{\hat{Z} - \varepsilon}.
\end{aligned}
$$

The last step follows from the fact that $\hat{Z} - \varepsilon \leq Z$ and thus $1/Z \leq 1/(\hat{Z} - \varepsilon)$ as long as $\hat{Z} - \varepsilon > 0$. $\qquad\square$