

Efficient Probabilistic Collision Detection for Non-Gaussian Noise Distributions

Jae Sung Park¹ and Dinesh Manocha²

Abstract—We present an efficient algorithm to compute tight upper bounds of collision probability between two objects with positional uncertainties, whose error distributions are represented with non-Gaussian forms. Our approach can handle noisy datasets from depth sensors, whose distributions may correspond to Truncated Gaussian, Weighted Samples, or Truncated Gaussian Mixture Model. We derive tight probability bounds for convex shapes and extend them to non-convex shapes using hierarchical representations. We highlight the benefits of our approach over prior probabilistic collision detection algorithms in terms of tighter bounds (10x) and improved running time (3x). Moreover, we use our tight bounds to design an efficient and accurate motion planning algorithm for a 7-DOF robot arm operating in tight scenarios with sensor and motion uncertainties.

I. INTRODUCTION

Efficient collision detection is an important problem in robot motion planning, physics-based simulation, and geometric applications. Earlier work in collision detection focused on fast algorithms for rigid convex polytopes and non-convex shapes and later extended to non-rigid models [1], [2]. Most of these methods assume that an exact geometric representation of the objects is known in terms of triangles or continuous surfaces [3] and the output of collision query is a simple binary outcome.

As robots navigate and interact with real-world objects, we need algorithms for motion planning and collision detection that can handle environmental uncertainty. In particular, robots operate with sensor data, and it is hard to obtain an exact shape or pose of an object. For example, depth cameras are widely used in robotics applications and the captured representations may have errors that correspond to lighting, calibration, or object surfaces [4]. This gives rise to probabilistic collision detection, where the goal is to compute the probability of in-collision state by modeling the uncertainty using some probabilistic distribution.

Many collision detection algorithms have been proposed to account for such uncertainties [5], [6], [7]. In practice, it is hard to analytically compute the collision probability for all probabilistic representations of uncertainties. Most prior work on probabilistic collision detection is limited to Gaussian distributions [8], [9]. However, these formulations may not work well when objects are captured by a depth sensor. Even the process of capturing static objects can result in




Fig. 1: We highlight the benefits of our novel probabilistic collision detection with a Truncated Gaussian error distribution. Our formulation is used to accurately predict the future human motion and integrated with a motion planner for the 7-DOF Fetch robot arm. As compared to prior probabilistic collision detection algorithms based on Gaussian distribution [9], our new method improves the running time by 2.5x and improves the accuracy of collision detection by 7.8x.

depth images, where the depth values can vary between consecutive frames due to different noise sources. The dynamic objects in the scene can pose additional problems due to sensor noise and the uncertainties introduced by the objects' motions. Moreover, Gaussian process dynamical models used to represent human motion [10] have an inherited uncertainty in the Gaussian variances as the central motion is represented using Gaussian means.

In many applications, it is necessary to use non-Gaussian models for uncertainties [11]. These include Truncated Gaussian with bounded domains for sensory noises [12], [13] to represent the position uncertainties for a point robot position [14]. Other techniques model the uncertainty as a Partially Observable Markov Decision Process (POMDP) [15], [16].

Main Results: We present efficient algorithms to compute the collision probability for error distributions corresponding to a variety of non-Gaussian models, including weighted samples and Truncated Gaussian (TG). Our approach is based on modeling the TG error distribution and represent the collision probability using a volume integral (Section 4). We present efficient techniques to evaluate the integral and highlight the benefits over prior methods for probabilistic collision detection. We evaluate their performance on synthetic as well as real-world datasets captured using depth cameras (Section 5). Furthermore, we show that our efficient probabilistic collision detection algorithm can be used for real-time robot motion planning of a 7-DOF manipulator in tight scenarios with depth sensors. Some novel components of our work include:

- A novel method to perform probabilistic collision detection for TG Mixture Models based on appropriately formulating of the vector field, and computing an up-

¹ Jae Sung Park with the Department of Computer Science, University of North Carolina at Chapel Hill, NC, USA jaesungp@cs.unc.edu

² Dinesh Manocha is with Department of Computer Science and Electrical & Computer Engineering, University of Maryland at College Park, MD, USA dmanocha@umd.edu

per bound using divergence theorem on the resulting integral. Moreover, we present an efficient method to evaluate this bound for convex and non-convex shapes.

- We show that TG outperforms normal Gaussian, and Truncated Gaussian Mixture Model (TGMM) outperforms Gaussian Mixture Model (GMM). In practice, probabilistic formulation is less conservative than prior methods and results in $5 - 9\times$ accuracy in terms of collision probability computation (Table 1).
- We have combined our probabilistic collision formulation with an optimization-based realtime robot motion planner that accounts for positional uncertainty from depth sensors. Our modified planner is less conservative in terms of computing paths in tight scenarios.

II. RELATED WORK

We give a brief overview of prior work on probabilistic collision detection.

A. Probabilistic Collision Detection for Gaussian Errors

Many approaches to compute the collision probability in uncertain robotic environments approximate the noises using a single Gaussian or a mixture of Gaussian distributions to simplify the computations. Such approaches are widely used in 2D environments for autonomous driving cars to avoid collisions with cars or pedestrians. Xu et al. [8] use Linear-Quadratic Gaussian to model the stochastic states of car positions on the road. Collision detection under uncertainty is performed by computing the Minkowski sum of Gaussian ellipse boundary and the rectangular car model and checking for overlap with other rectangular car model. Park et al. present an efficient algorithm to compute an upper bound of the collision probability with Gaussian error distributions [9]. This approach can be extended to Truncated Gaussian because the probability density function (PDF) of a Truncated Gaussian inside its ellipsoidal domain has the same value as that of the PDF of a Gaussian. Therefore, the upper bound computed using [9] also holds for Truncated Gaussian error distributions, but the bound is not tight. Moreover, a Truncated Gaussian distribution has a bounded ellipsoidal domain and the integral computations outside the domain can be omitted. As compared to this approach, our new algorithm improves the upper bound and the running time, as shown in Section 5.

B. Probabilistic Collision Detection for Non-Gaussian Errors

The collision probability for non-Gaussian error distributions can be computed with Monte Carlo sampling [17]. However, these methods are much slower ($10 - 1000$ times), as compared to probabilistic algorithms that use Gaussian forms of error distributions [9]. Althoff et al. [18] use a non-Gaussian probability distribution model on the future states of other cars on the road, based on their positions, speeds, and road geometry. They use a 2D grid discretization of the state space and Markov chain to compute the probability that a car belongs to a cell. This method assumes that the

environment sensors has no noise. Lambert et al. [19] use a Monte Carlo approach, taking advantage of the probability density function represented as a Gaussian. Other methods have been proposed for point clouds using classification [7] or Monte-Carlo integration [20]. Approaches based on Partially Observable Markov Decision Processes (POMDPs) make efficient decisions about the robot actions in a partially observable state in an uncertain environment [11], [21]. Some applications using POMDPs [22] have been developed to avoid collisions in an uncertain environment, where the uncertainty is represented with a non-Gaussian probability distribution. Our approach for non-Gaussian distributions is different and complementary with respect to these methods.

C. Probabilistic Collision Detection: Applications

Many approaches have been proposed for collision checking for general applications. Aoude et al. [23] represent the uncertainty model for point obstacles as a Gaussian Process and positional error is represented by a Gaussian distribution that propagates over a discretized time domain. The upper bound on the collision probability is computed on the Gaussian positional error with an $\text{erf}(\cdot)$ function for a point obstacle. Fisac et al. [24] compute the collision probability between the dynamic human motion and a robot, and use that value for robot motion planning in the 3D workspace. This algorithm models the human motion based on human dynamics, discretizes the 3D workspace into smaller grids, and integrates the cell probabilities over the volume occupied by the robot. Probabilistic collision detection for a Gaussian error distribution [9] has been used for optimization-based robot motion planning. The collision constraint used in the optimization formulation is that the collision probability should be less than 5% at any robot configuration in the resulting trajectory. However, with Gaussian error distributions, the upper bound of collision probability is rather conservative. As a result, these approaches do not work well in tight spaces or narrow passages.

III. OVERVIEW

In this section, we introduce the terminology used in the paper and give an overview of our approach. Our algorithm is designed for environments, where the scene data is captured using sensors and only partial observations are available. In this case, the goal is to compute the *collision probability* between two objects, when one or both objects are represented with uncertainties and some of the input information such as positions or orientations of polygons or point clouds are given as probability distributions

A. Probabilistic Collision Detection

The input of the probabilistic collision detection is two 3D shapes A and B , and two 3D positional error distributions P_A and P_B that are probabilistically independent of each other. The positional error distributions P_A and P_B denote the probability density function over 3D space of translations from the origins of objects A and B , respectively. The output of the algorithm is p_{col} , the probability of in-collision state

between A and B , where the objects can be translated with the error distributions.

The collision probability p_{col} , given two input shapes A and B and the error distributions p_A and p_B , can be formulated as

$$p_{col} = \iiint_{\epsilon_A} \iiint_{\epsilon_B} I((A + \epsilon_A) \cap (B + \epsilon_B) \neq \emptyset) p(\epsilon_A) p(\epsilon_B) d\epsilon_A d\epsilon_B, \quad (1)$$

$$\epsilon_A \sim P_A, \quad \epsilon_B \sim P_B, \quad (2)$$

where $I(\cdot)$ is an indicator function which yields 1 if the condition is true and 0 otherwise, and ϵ_A and ϵ_B are the displacement vectors for A and B with the probability distribution P_A and P_B .

To generalize, we shift only one object A by $\epsilon = \epsilon_A - \epsilon_B$ which follows a probabilistic distribution P_{AB} , instead of shifting the two objects separately by ϵ_A and ϵ_B . Because of the independence of probabilistic distributions P_A and P_B , the convolution P_{AB} of P_A and P_B can be expressed as:

$$f_{AB}(\mathbf{x}) = \iiint_{\mathbf{y}} f_A(\mathbf{y}) f_B(\mathbf{x} - \mathbf{y}) d\mathbf{y}, \quad (3)$$

where f_{AB} , f_A , f_B are the probability density functions of P_{AB} , P_A , P_B , respectively.

B. Probabilistic Collision Detection for Gaussian Error

The general probabilistic collision detection problem is hard to solve, when the error distributions P_A and P_B have any arbitrary form. The convolution operator in (Equation (3)) can be hard to formulate in the general case. However, it is known that the convolution of two Gaussians is also Gaussian. This generalizes the use of two error distributions into one, yielding the following:

$$p_{col} = \iiint_{\epsilon} I((A + \epsilon) \cap B) \neq \emptyset p(\epsilon) d\epsilon \quad (4)$$

$$= \iiint_{\epsilon} I(\epsilon \in (-A) \oplus B) p(\epsilon) d\epsilon, \quad \epsilon \sim P_{AB}, \quad (5)$$

where \oplus denotes the Minkowski sum operator between two shapes. Probabilistic collision detection with the Gaussian distribution condition can be solved efficiently [9], where P_A and P_B also correspond to Gaussian distributions. This algorithm computes a good upper bound on collision probability for convex and non-convex shapes by efficiently linearizing the Gaussian along the minimum displacement vector direction. In practice, the resulting bounds are conservative.

IV. TRUNCATED GAUSSIAN MIXTURE MODEL ERROR DISTRIBUTION

In this section, we present an efficient algorithm for Truncated Gaussian Mixture Model (TGMM) error distributions, which is a more general type of noise model for robotics applications. To compute the collision probability for TGMM, we first introduce the solutions for simpler error distributions corresponding to Truncated Gaussian (TG) and Weighted Samples (WS). We combine these two algorithms to design an algorithm for a multiple Truncated Gaussian error distribution model.

A. Truncated Gaussian Mixture Models

A TGMM consists of multiple Truncated Gaussian (TG) distributions, each distribution with a truncated domain. The probability density function of a TG, f_{TG} , can be formulated as:

$$f_{TG}(\mathbf{x}; \mu, \Sigma, r) = \begin{cases} \frac{1}{\eta} g(\mathbf{x}; \mu, \Sigma) & (\mathbf{x} - \mu)^T \Sigma^{-1} (\mathbf{x} - \mu) \leq r \\ 0 & \text{otherwise} \end{cases}, \quad (6)$$

where g is the probability density function of a Gaussian, μ is the mean, Σ is the variance, r is the radius of bound in the coordinates of the principal axes, and η is the truncation rate used to compensate the loss of truncated volume of probability outside the bound. A TGMM consists of n TGs with multiple weights w_i . The probability density function of the TGMM, f_{TGMM} , can be formulated using the definition of f_{TG} in Equation (6), as:

$$f_{TGMM}(\mathbf{x}) = \sum_{i=1}^n w_i f_{TG}(\mathbf{x}; \mu_i, \Sigma_i, r_i), \quad \sum_{i=1}^n w_i = 1. \quad (7)$$

As the radii of TGs decrease and converge to zero, the probability model behaves like a discrete probability distribution, which we call Weighted Samples (WS). The WS is a discrete probability distribution, formulated as:

$$P(\mathbf{X} = \mathbf{x}_i) = w_i, \quad \sum_{i=1}^n w_i = 1 \quad (8)$$

where \mathbf{x}_i is a sample in \mathbb{R}^d , and w_i is a weight of the sample, for $i = 1, \dots, n$.

B. Collision Probability for Truncated Gaussian

The TG is a Gaussian with a specific form of bounded domain. The bounded domain for 3D Truncated Gaussians is an ellipsoid, centered at the Gaussian mean and having the same principal axes as those of Gaussian variances. The TG is formulated with a collision probability function as

$$p_{col} = \iiint_{V_{AB}} f_{TG}(\mathbf{x}; \mu, \Sigma, r) d\mathbf{x}, \quad (9)$$

where $V_{AB} = -A \oplus B$, f_{TG} is the probability density function for TG, μ is the mean, Σ is the variance, r is the radius of bound in the coordinates of the principal axes, and η is the normalization constant used to compensate the loss of truncated volume of probability outside the bound. Because f_{TG} has the value of a Gaussian multiplied by η inside the boundary, the integral volume becomes $-A \oplus B \cap V_{TG}$, where V_{TG} is the valid volume of Truncated Gaussian. The collision probability corresponds to

$$p_{col} = \frac{1}{\eta} \iiint_{V_{AB} \cap V_{TG}} f_{TG}(\mathbf{x}; \mu, \Sigma, r) d\mathbf{x}. \quad (10)$$

The TG has its center at μ and principal axes with different lengths determined by Σ . To normalize the function, a

transformation $T = \Sigma^{-1/2} - \mu I$ is applied to the coordinate system, which changes Equation (10) to

$$p_{col} = \frac{1}{\eta \det \Sigma} \iiint_{V'_{AB} \cap V'_{TG}} f_{TG}(\mathbf{x}; \mathbf{0}, I, r) d\mathbf{x}, \quad (11)$$

$$V'_{AB} = T(V_{AB}), \quad V'_{TG} = T(V_{TG}). \quad (12)$$

In the transformed coordinate system, V'_{TG} is a sphere of radius r .

Unfortunately, there is no explicit or analytic form of solution for the integral of a Gaussian distribution over the intersection of a non-convex volume V'_{AB} and a ball V'_{TG} . In order to simplify the problem, we initially assume that A and B are convex, and so are V_{AB} and V'_{AB} . Instead of computing the exact integral, we compute an upper bound on the collision probability. The computation of collision probability reduces to the computation of the integral

$$\iiint_{V'} g(\mathbf{x}; \mathbf{0}, I) d\mathbf{x}, \quad (13)$$

where $V' = V'_{AB} \cap V'_{TG}$, and $g(\cdot)$ is the Gaussian probability density function.

From the convexity of V'_{AB} , the minimum distance vector \mathbf{d}' between the origin and V'_{AB} can be computed by using the GJK algorithm [25] between A' and B' , which are transformed from A and B by T . Let \mathbf{n}'_d be the unit directional vector of \mathbf{d}' . Then, by the Cauchy-Schwarz inequality $(\mathbf{x} \cdot \mathbf{n}'_d)^2 \leq \|\mathbf{x}\|^2$, the integral is bounded by

$$p_{col} \leq \iiint_{V'} \frac{1}{\sqrt{8\pi^3}} \exp\left(-\frac{1}{2}(\mathbf{x} \cdot \mathbf{n}'_d)^2\right) d\mathbf{x}. \quad (14)$$

The integrand of the upper bound term behaves as a 1D Gaussian function instead of being the 3D function. We use the divergence theorem to compute the upper bound on collision probability (14).

$$\iiint_{V'} \operatorname{div}(\mathbf{F}) dV = \iint_{S'} (\mathbf{F} \cdot \mathbf{n}_S) dS, \quad (15)$$

where \mathbf{F} is a vector field, S' is the surface of V' , dS is an infinitesimal area for integration, and \mathbf{n}_S is the normal vector of dS . This converts the volume integral to a surface integral. Let's define \mathbf{F} as

$$\mathbf{F}(\mathbf{x}) = \frac{1}{2\pi} \left(1 + \operatorname{erf}\left(\frac{\mathbf{x} \cdot \mathbf{n}'_d}{\sqrt{2}}\right) \right) \mathbf{n}'_d, \quad (16)$$

where $\operatorname{erf}(\cdot)$ is the 1D Gaussian error function. Note that \mathbf{F} is a vector field with a single direction \mathbf{n}'_d . The directional derivative of $\mathbf{F}(\mathbf{x})$ along any directional vector orthogonal to \mathbf{n}'_d is zero because \mathbf{F} varies only along \mathbf{n}'_d . The divergence of \mathbf{F} thus becomes $(\partial \mathbf{F} / \partial \mathbf{n}'_d)$, and this is equal to the function in Equation (14).

We apply the divergence theorem in Equation (15) to the volume integral on V' in Equation (14). Note that V' is a 3D volume intersection between a non-convex polytope V'_{AB} and a ball V'_{TG} . The surface integral on the intersection between




Fig. 2: (a) Contour plots of the bivariate TG distribution. (b) Contour plots of the bounded function \mathbf{F} for TG are not used in the calculation of collision probability and thereby reduce the running time of collision probability computation.

a non-convex polytope and a ball can be decomposed into two parts and bounded by the sum of two components as

$$\sum_i \iint_{\Delta S'_i} (\mathbf{F} \cdot \mathbf{n}'_i) dS + \iint_{S'_{TG}} (\mathbf{F} \cdot \mathbf{n}_S) dS, \quad (17)$$

where S'_i is the i -th triangle of V'_{AB} inside V'_{TG} , \mathbf{n}'_i is the normal vector of $\Delta S'_i$, and S'_{TG} is the spherical boundary of V'_{TG} outside of a plane defined by \mathbf{d}' . The second term corresponds to the spherical domain of the normalized Truncated Gaussian with the truncation rate η . The magnitude of F on the spherical boundary V'_{TG} is upperly bounded by $(1 - \eta)$, because it is the cumulative distribution function on the boundary. The surface area of S'_{TG} is less than $\pi \|\mathbf{d}'\|^2$. This can be used to express a bound based on the following lemma.

Lemma 4.1: The collision probability represented in a volume integral is upperly bounded by a surface integral as follows:

$$p_{col} = \iiint_{V'} g(\mathbf{x}; \mathbf{0}, I) d\mathbf{x} \quad (18)$$

$$\leq \sum_i \iint_{\Delta S'_i} (\mathbf{F} \cdot \mathbf{n}_i) dS + \pi(1 - \eta) \|\mathbf{d}'\|^2, \quad (19)$$

where F is a vector field in 3D space whose maximum magnitude is $1/\pi$, and S_i is the i -th triangle of V'_{AB} that is inside V'_{TG} .

Because the error function integral over a triangle domain is hard to compute, the upper bound on the integral is evaluated as

$$\sum_i \iint_{\Delta S'_i} (\mathbf{F} \cdot \mathbf{n}_i) dS \quad (20)$$

$$\leq \sum_i \left(\max_{j=1,2,3} \mathbf{F}(S_{ij}) \cdot \mathbf{n}_i \right) \operatorname{Area}(\Delta S'_i), \quad (21)$$

where S_{ij} is the j -th vertex of the triangle S_i for $j \in \{1, 2, 3\}$. The upper bound on the collision probability corresponds to the sum of the maximum of F at the points of each triangle, multiplied by the area of the triangle, over the surface of $V'_{AB} \cap V'_{TG}$.

C. Efficient Evaluation of the Integral

In order to reduce the running time of computing the surface integral, we take advantage of the bound of TG. The domain of surface integral is $V' = V'_{AB} \cap V'_{TG}$, where




Fig. 3: The upper bound of collision probability with uncertainty approximated as Gaussian, Gaussian Mixture, and Weight Samples. The X-axis is the true collision probability computed using Monte Carlo methods, and Y-axis is the computed probability using different methods. The computed upper bound for Gaussian Mixture and Weights Samples are closer to the ground truth/exact answer, than that for a single Gaussian approximation. The collision probability over-estimation with TGMM is reduced by 90%, compared to the one with Gaussian distribution.

V'_{AB} consists of many triangles and V'_{TG} is a sphere of radius r . This sphere is tightly bounded by a cube, with one normal parallel to the direction of shortest displacement vector \mathbf{d}' . Therefore, the triangles of V'_{AB} that are outside the cube do not count towards the surface integral. So, we accumulate the upper bound function value in Equation (21) only for the triangles that lie inside the cube boundary, and ignore the triangles outside the boundary. For the triangles that intersect the cube boundary, the upper bound function value is computed for the intersecting primitives. Because the approximated integral for collision probability is bounded by the cube, primitives outside the cube can be ignored in terms of calculating the upper bound of collision probability. Limiting the computation to the truncated primitives can accelerate the running time.

In order to perform this computation for non-convex primitives, we construct bounding volume hierarchies (BVHs) for A and B , with each bounding volume being an oriented bounding box. During the traversal of the BVHs, the oriented bounding boxes are first transformed by T . The transformed bounding volumes are still convex primitives, the surface integral can be obtained using Equation (21).

D. Error Distribution as Weighted Samples

For the weighted samples, the probability distributions are given by multiple points p_i with weights w_i , yielding a discrete probability distribution, as described in Equation (8). The collision probability of Equation (4-5) for the weighted samples is given as:

$$p_{col} = \sum_{i=1}^n w_i I(\mathbf{x}_i \in (-A) \oplus B), \quad \sum_{i=1}^n w_i = 1, \quad (22)$$

where w_i is weight and $I(\cdot)$ is the indicator function which yields 1 if the statement inside is true or 0 otherwise. The formulation is the weighted average of n collision detection results. A simple solution to this problem is to run exact collision detection algorithms n times and sum up the weights of in-collision cases. However, this results in an $O(n)$ and we use BVHs to accelerate that computation.

We have the bounding volumes for the weighted samples and the two polyhedra. When there is no overlap between the bounding boxes, it implies that there is no collisions between two shapes for all weighted samples in the corresponding bounding volume. If the bounding volumes overlap, there may be a collision for each weighted samples, and the bounding volumes of the children are checked recursively for collisions. Each of these bounding volume checks can be performed in $O(1)$ time.

If we want to compute an upper bound of collision probability, the running time can be further reduced by replacing detailed computation of collision probability with a simple upper bound. We introduce a user-defined parameter δ which we call the “confidence level”. During the traversal of bounding volume traversal tree, the upper bound of collision is the sum of weights of samples that belong to the bounding volume. If the upper bound is less than the confidence level $1 - \delta$, the traversal stops and the sub-routine returns the sum of weights as an upper bound of collision probability.

In order to reduce the time complexity for more complex forms of error distributions, we construct a Bounding Volume Hierarchy (BVH) [1] over the error distributions of mixture models with Oriented Bounding Boxes (OBBs) [26] and apply the collision probability algorithm on its nodes, which are convex primitives. We construct a BVH for the weighted samples and for Truncated Gaussian Mixture Models in $O(n)$ time complexity. The BVH is generated from the root node that contains every Truncated Gaussians, and the bounding volume for the root node is computed by minimizing the volume of the oriented bounding box. Next, the bounding volume is split at the center along the longest edge and two child BVH nodes are generated, each containing appropriate samples. This process is repeated till the leaf nodes.

E. Error Distribution as Truncated Gaussian Mixture Models

For TGMM the probability distribution is given as:

$$p_{col} = \sum_{i=1}^n w_i \iiint_{V_{AB}} \eta_i f_{TGMM}(\mathbf{x}; \mu_i, \Sigma_i, r_i) d\mathbf{x}, \quad (23)$$

where n is the number of Truncated Gaussians (TGs), w_i is the weight of each TG, and μ_i , Σ_i and r_i are the mean, variance and radius of TGs, respectively. The overall algorithm for TGMM is obtained by combining the two previous algorithms. A change from the algorithm for weighted samples is that the BVH is constructed for n TGs with their ellipsoid bounds instead of the point samples. The details of algorithms and pseudo-codes are given in the appendix [27].

V. PERFORMANCE AND ANALYSIS

In this section, we describe our implementation and highlight the performance of our probabilistic collision detection algorithms on synthetic and real-world benchmarks. Furthermore, we measure the upper bound of collision probabilities and speedups for algorithms with different noise distributions, compared to the exact collision probability computed by Monte Carlo method.

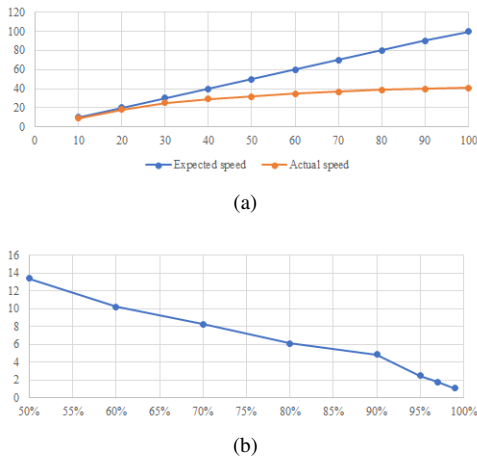


Fig. 4: (a) Speedup of weighted samples (expected) case compared to Monte Carlo (actual) with between 10 to 100 samples (X-axis). (b) Speedup of Truncated Gaussian case, compared to the running time of probabilistic collision detection with a Gaussian. X-axis is the untruncated volume of Gaussian, meaning 100% is the Gaussian and lower value indicates smaller bound. As the truncation boundary shrinks up to 50% of the volume of Gaussian, the algorithm with TG is 14x faster times than the algorithm Gaussian distribution.

A. Probabilistic Collision Detection: Performance

For the translational error distribution, we first generate a ground truth distribution by randomizing the parameters of TGMM. Next, we sample 100,000 points from the distribution. We run expectation-maximization from these samples to find parameters of single Gaussian, single TG, WS, and TGMM. The resultant distributions are different from the ground truth and count toward collision probability over-estimation.

Figure 3 shows the collision probabilities of noise models approximated with Gaussian, TGMM of 10 TG distributions, and 100 WS. We observe that the approximation with a single Gaussian yields rather high and conservative value of collision probability, compared to the ground truth collision probability. Figure 4 (a) shows the speedup of probabilistic collision detection with WS over the Monte Carlo method. The collision probability computation with Monte Carlo counts the sum of sample weights for every Weighted Sample that is in collision, and is similar to an exact collision detection algorithm. Thus, the running time of Monte Carlo increases linearly as the number of Weighted Samples increases. Figure 4 (b) shows the speedup of probabilistic collision detection with a TG model compared to probabilistic collision detection with a Gaussian error distribution. In case of TGs, BVH traversal is not performed when the truncation boundary of TGs does not overlap with the Minkowski sum of bounding volumes for the two objects. On the other hand, for Gaussian distributions there are no truncation boundaries and the BVH traversal continues. Therefore, we observe a speedup with TGs over Gaussian, as shown in Figure 1 and Table 1. Overall, the speedup depends on the range of truncation. A smaller truncation boundary




Fig. 5: (a) A captured RGBD image. The depth values of the table and the wood block have noises, even in adjacent frames. The TG noise of each point particle of the wood block contributes to the overall TGMM model. (b) A reconstructed 3D model of a wood block with TGMM, which is bounded around the wooden block and is more accurate than the Gaussian distribution, which has an unbounded probability density function. (c) A reconstructed 3D robot environment with error distributions on the table and the wood blocks. The wood blocks placed in a zig-zag pattern result in 4 narrow passages for the robot.

results in faster performance of our probabilistic collision detection algorithm.

B. Sensor Noise Models for Static Obstacles

In a real-world setting, we add a noise model to the point cloud data. The variance is chosen based on the Kinect sensor uncertainty. The input depth images have noise in each pixel and, according to [4], the noise of each pixel can be approximated with a 1D Gaussian. Thus, noise in the pixels of an object are combined with a Gaussian Mixture noise model for the object. Depth images of the wooden blocks and the table have noise in each pixel. We capture sequences of the depth images. In the experiments, the approximate poses of tables and wooden blocks are known a priori. On the boundary of the wooden blocks, some pixels are always classified as a wooden block, some other pixels are classified as a wooden block or as background in different frames, and other pixels are always classified as background. From these boundary pixels, the variance in x- and y-axis noise on the Kinect sensor coordinate frame can be set to the thickness of the boundary. The variance in z-axis motion is computed from the always-wooden-block pixels. The depth values on those pixels differ from frame to frame. After the variance of noise Gaussian is calculated, the mean and variance of a Truncated Gaussian are the same. In our benchmarks, the truncation rate η is such that the integral in the ellipsoidal boundary is 90%. With the fixed truncation rate, the positional error is bounded around the wooden blocks, unlike the positional error represented by Gaussians with an unbound domain. The confidence level δ is related to the robot motion planner, constraining that the collision probability between the robot and the objects should be less than $1 - \delta$ at any robot trajectory point. In our benchmarks, δ is set to 95%.

Figure 5 shows a captured depth image, a noise distribution modeled using Truncated Gaussian Mixture Model, and a reconstructed 3D environment with noises. The pixels on the boundary of the object have higher variance in terms of

noise. So, the Truncated Gaussian Mixture noise model may have some Gaussians with higher variance. A principal axis for those boundary pixels is perpendicular to the boundary direction.

C. Robot Motion Planning

The probabilistic collision detection algorithm is used in optimization-based motion planning [28]. We use a 7-DOF Fetch robot arm in the motion planning. We highlight its performance in terms of improved accuracy and faster running time in Figure 1. In the robot environment, there is a table in front of the robot and the wood blocks on the table are the static obstacles of the environment. The wood blocks are placed in a manner that pairs of them result in a narrow passage for the robot’s end-effector. The environment is captured using two depth sensors. One sensor is the Primesense Carmine 1.09 sensor, the robot head camera. Another one is the Microsoft Kinect 2.0 sensor installed in the opposite direction with respect to the robot. The point clouds of the table and the wood blocks captured by the two sensors are used to reconstruct the environment. In this case, the reconstructed table surface and wood block obstacles have errors due to the noise in the depth sensors. Figure 5 (c) shows the reconstructed environment from depths sensors and the error distributions around the obstacles. The robot arm’s task is to move a wood block, drawing a zig-zag pattern that passes through the narrow passages between the wood blocks. The objective is to compute a robot trajectory that minimizes the distance between the robot’s end-effector and the table, without resulting in any collisions. The following metrics are used to evaluate the performance:

- **Collision Probability Over-estimation:** Because we compute an upper bound of the collision probability in our algorithm, we measure the extent of collision probability over-estimation, the gap between the upper bound of collision probability and the actual collision probability.
- **Running Time:** The running time of the collision detection algorithm. This excludes the running time of motion planning algorithm.
- **# Passages:** The successful number of passes the robot makes between wood blocks.
- **Success Rate:** The ratio of collision-free trajectories among the total number of trajectory executions. For each execution, the wood block’s positions are set following the error distribution.
- **Distance:** The distance between the robot’s end-effector and the table. A lower value is better, as it implies that the robot can interact with the environment in close proximity.

The collision probability over-estimation, running time, distance values are measured for robot poses of every 1/30 seconds over the robot trajectories.

We compare the performances of 5 different collision detection algorithms: exact collision detection with static obstacles without environment uncertainties (CD-Obstacles), exact collision detection with point clouds (CD-Points),

probabilistic collision detection with Gaussian errors (PCD-Gaussian) [9], PCD with Truncated Gaussian (PCD-TG), PCD with weighted samples (PCD-WS), and PCD with Truncated Gaussian Mixture Model (PCD-TGMM). The weighted samples are drawn from the TGMMs. For each TG distribution, one sample is drawn from its center. Other samples are drawn from three icosahedrons with the same centers and different radii by uniformly dividing the truncation radius.

Benefits of Truncated Gaussian: Table I shows the results of robot motion planning in the various scenarios with different algorithms. The robot motion planning without uncertainties (CD-Obstacles) operates perfectly. However, under sensor uncertainties, the exact collision detection with the point clouds (CD-Points) works poorly. Due to the sensor uncertainties, a high error at a pixel affects the collision detection query accuracy and the performance in narrow passages and success rate. Our probabilistic collision detection algorithms results in better performances than exact collision detection algorithm under the environment with sensor uncertainties. Compared to PCD-Gaussian and PCD-GMM, our algorithms for non-Gaussian distributions (PCD-WS, PCD-TGMM) demonstrate better performances w.r.t. different metrics.

Human Motion Prediction: We also evaluate our algorithm on scenarios with humans operating close to the robot, as shown in Fig. 1 and the video. In order to handle the uncertainty of future human motion, we use probabilistic collision detection between the robot and the predicted future human pose. Improvement in the accuracy of motion prediction results in better trajectories in terms of being collision-free, smoother and being able to handle tight scenarios [29]. We highlight the benefits of our PCD-TGMM algorithm on the resulting trajectory computation in the video.

VI. CONCLUSION AND LIMITATIONS

We present efficient probabilistic collision detection algorithms for the following forms of non-Gaussian error distributions: Truncated Gaussian, Weighted Samples, and Truncated Gaussian Mixture Model. Compared to the exact collision detection algorithm and prior probabilistic collision detection algorithms for Gaussian error distribution, our new method can compute a tighter upper bound of collision probability and improves the running time. We have integrated this algorithm with a motion planner and highlights its benefits in narrow passage scenarios with a 7-DOF robot arm. Our algorithm can be used to model non-Gaussian error distributions from noisy depth sensors and predicted human motion models.

Our approach has some limitations. The truncation on the Truncated Gaussians has a form of ellipsoid with the center and principal axes that is the same as the original Gaussian distribution. The ellipsoidal shape of the truncation boundary may not be sufficient for representing general error distributions. Another realistic possibility would be to truncate using planar boundaries. In our future work, we would like to develop an algorithm for more generic non-Gaussian positional errors. We only consider the positional errors

Algorithm	Collision Probability Over-estimation (%p)			Running Time (ms)			# Passages	Success Rate	Distance (cm)
	Min	Max	Avg	Min	Max	Avg			
CD-Obstacles	0	0	0 (0)	2.8	8.0	3.5 (0.75)	4/4	10/10	1.2 (0.09)
CD-Points	0	100	23 (16)	2300	2800	2500 (130)	1/4	2/10	13 (8.8)
PCD-Gaussian [9]	8.6	35	15 (5.2)	15	100	27 (23)	3/4	8/10	7.9 (2.6)
PCD-TG	2.5	13	5.0 (3.3)	5.2	86	12 (4.6)	3/4	7/10	6.0 (1.3)
PCD-WS	0.72	8.1	1.5 (0.60)	130	220	180 (32)	4/4	9/10	4.5 (0.66)
PCD-GMM [9]	1.9	7.7	6.2 (3.3)	85	480	240 (55)	3/4	8/10	7.1 (1.0)
PCD-TGMM	0.26	3.3	0.8 (0.32)	35	130	97 (17)	4/4	9/10	3.7 (0.46)

TABLE I: **Performance of probabilistic collision detection algorithms:** Evaluated as part of a motion planner with sensor data. The collision probability over-estimation is shown as percent point (%p) with the minimum and maximum over-estimation. The values corresponds to the average over the time of the robot trajectory with standard deviation in parenthesis. The best performance is obtained PCD-TGMM algorithm in terms of collision probability estimation (i.e. tight bounds), successful handling of narrow passages, computing collision-free trajectories, among different algorithms.

on obstacles and omit the rotational errors. The rotational error cannot be approximated by Truncated Gaussian Mixture Model. As part of future work, we would like to represent a rotational error distribution in the quaternion space, or in the affine space. Furthermore, we assume that the noises of two objects A and B are independent, even though they may arise from the same source.

ACKNOWLEDGMENTS

This research is supported in part by ARO grant W911NF18-1-0313 and Intel.

REFERENCES

- [1] M. C. Lin and D. Manocha, "Collision and proximity queries," in *Handbook of Discrete and Computational Geometry: Collision detection*. CRC Press, 2003, pp. 787–808.
- [2] J. T. Klosowski, M. Held, J. S. Mitchell, H. Sowizral, and K. Zikan, "Efficient collision detection using bounding volume hierarchies of k-dops," *IEEE transactions on Visualization and Computer Graphics*, vol. 4, no. 1, pp. 21–36, 1998.
- [3] A. Greß, M. Guthe, and R. Klein, "Gpu-based collision detection for deformable parameterized surfaces," in *Computer Graphics Forum*, vol. 25, no. 3. Wiley Online Library, 2006, pp. 497–506.
- [4] K. Khoshelham and S. O. Elberink, "Accuracy and resolution of Kinect depth data for indoor mapping applications," *Sensors*, vol. 12, no. 2, pp. 1437–1454, 2012.
- [5] R. B. Rusu, I. Alexandru, B. Gerkey, S. Chitta, M. Beetz, L. E. Kavraki et al., "Real-time perception-guided motion planning for a personal robot," in *Proceedings of IROS*. IEEE, 2009, pp. 4245–4252.
- [6] K.-H. Bae, D. Belton, and D. D. Lichti, "A closed-form expression of the positional uncertainty for 3d point clouds," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 31, no. 4, pp. 577–590, 2009.
- [7] J. Pan, S. Chitta, and D. Manocha, "Probabilistic collision detection between noisy point clouds using robust classification," in *International Symposium on Robotics Research (ISRR)*, 2011.
- [8] W. Xu, J. Pan, J. Wei, and J. M. Dolan, "Motion planning under uncertainty for on-road autonomous driving," in *ICRA*. IEEE, 2014, pp. 2507–2512.
- [9] J. S. Park, C. Park, and D. Manocha, "Efficient probabilistic collision detection for non-convex shapes," in *Proceedings of ICRA*. IEEE, 2017, pp. 1944–1951.
- [10] J. M. Wang, D. J. Fleet, and A. Hertzmann, "Gaussian process dynamical models for human motion," *IEEE transactions on pattern analysis and machine intelligence*, vol. 30, no. 2, pp. 283–298, 2008.
- [11] H. Kurniawati and V. Yadav, "An online POMDP solver for uncertainty planning in dynamic environment." ISRR, 2013.
- [12] N. L. Johnson, S. Kotz, and N. Balakrishnan, "Lognormal distributions," *Continuous univariate distributions*, vol. 1, pp. 207–227, 1994.
- [13] D. Simon, *Optimal state estimation: Kalman, H infinity, and nonlinear approaches*. John Wiley & Sons, 2006.
- [14] S. Patil, J. Van Den Berg, and R. Alterovitz, "Estimating probability of collision for safe motion planning under gaussian motion and sensing uncertainty," in *Proceedings of ICRA*. IEEE, 2012, pp. 3238–3244.
- [15] M. Rafieisakhaei, S. Chakravorty, and P. Kumar, "Non-gaussian slap: Simultaneous localization and planning under non-gaussian uncertainty in static and dynamic environments," *arXiv preprint arXiv:1605.01776*, 2016.
- [16] K. M. Seiler, H. Kurniawati, and S. P. Singh, "An online and approximate solver for POMDPs with continuous action space," in *Proceedings of ICRA*. IEEE, 2015, pp. 2290–2297.
- [17] N. E. Du Toit and J. W. Burdick, "Probabilistic collision checking with chance constraints," *Robotics, IEEE Transactions on*, vol. 27, no. 4, pp. 809–815, 2011.
- [18] M. Althoff, O. Stursberg, and M. Buss, "Model-based probabilistic collision detection in autonomous driving," *IEEE Transactions on Intelligent Transportation Systems*, vol. 10, no. 2, pp. 299–310, 2009.
- [19] A. Lambert, D. Gruyer, and G. S. Pierre, "A fast monte carlo algorithm for collision probability estimation," in *Proceedings of ICARCV*. IEEE, 2008, pp. 406–411.
- [20] C. Park, J. S. Park, and D. Manocha, "Fast and bounded probabilistic collision detection in dynamic environments for high-dof trajectory planning," *Proceedings of WAFR*, 2016.
- [21] H. Bai, S. Cai, N. Ye, D. Hsu, and W. S. Lee, "Intention-aware online POMDP planning for autonomous driving in a crowd," in *Proceedings of ICRA*. IEEE, 2015, pp. 454–460.
- [22] J. Van den Berg, D. Wilkie, S. J. Guy, M. Niethammer, and D. Manocha, "LQG-Obstacles: Feedback control with collision avoidance for mobile robots with motion and sensing uncertainty," in *Proceedings of ICRA*. IEEE, 2012, pp. 346–353.
- [23] G. S. Aoude, B. D. Luders, J. M. Joseph, N. Roy, and J. P. How, "Probabilistically safe motion planning to avoid dynamic obstacles with uncertain motion patterns," *Autonomous Robots*, vol. 35, no. 1, pp. 51–76, 2013.
- [24] J. F. Fisac, A. Bajcsy, S. L. Herbert, D. Fridovich-Keil, S. Wang, C. J. Tomlin, and A. D. Dragan, "Probabilistically safe robot planning with confidence-based human predictions," *arXiv preprint arXiv:1806.00109*, 2018.
- [25] E. G. Gilbert, D. W. Johnson, and S. S. Keerthi, "A fast procedure for computing the distance between complex objects in three-dimensional space," *IEEE Journal on Robotics and Automation*, vol. 4, no. 2, pp. 193–203, 1988.
- [26] S. Gottschalk, M. C. Lin, and D. Manocha, "Obbtrees: A hierarchical structure for rapid interference detection," in *Proceedings of SIG-GRAPH*. ACM, 1996, pp. 171–180.
- [27] J. S. Park and D. Manocha, "Efficient probabilistic collision detection for non-gaussian noise distributions," *arXiv preprint arXiv : 1902.10252*, 2019.
- [28] C. Park, J. Pan, and D. Manocha, "ITOMP: Incremental trajectory optimization for real-time replanning in dynamic environments," in *Proceedings of ICAPS*, 2012.
- [29] J. S. Park, C. Park, and D. Manocha, "I-planner: Intention-aware motion planning using learning-based human motion prediction," *The International Journal of Robotics Research*, vol. 38, no. 1, pp. 23–39, 2019.