Leveraging Distributional Bias For Reactive Collision Avoidance under Uncertainty: A Kernel Embedding Approach

Anish Gupta[1], Arun Kumar Singh[2], and K. Madhava Krishna[1]

***Abstract*— Many commodity sensors that measure the robot and dynamic obstacle's state have non-Gaussian noise characteristics. Yet, many current approaches treat the underlying uncertainty in motion and perception as Gaussian, primarily to ensure computational tractability. On the other hand, existing planners working with non-Gaussian uncertainty do not shed light on leveraging distributional characteristics of motion and perception noise, such as bias for efficient collision avoidance.**

This paper fills this gap by interpreting reactive collision avoidance as a distribution matching problem between the collision constraint violations and Dirac Delta distribution. To ensure fast reactivity in the planner, we embed each distribution in Reproducing Kernel Hilbert Space and reformulate the distribution matching as minimizing the Maximum Mean Discrepancy (MMD) between the two distributions. We show that evaluating the MMD for a given control input boils down to just matrix-matrix products. We leverage this insight to develop a simple control sampling approach for reactive collision avoidance with dynamic and uncertain obstacles.

We advance the state-of-the-art in two respects. First, we conduct an extensive empirical study to show that our planner can infer distributional bias from sample-level information. Consequently, it uses this insight to guide the robot to good homotopy. We also highlight how a Gaussian approximation of the underlying uncertainty can lose the bias estimate and guide the robot to unfavorable states with a high collision probability. Second, we show tangible comparative advantages of the proposed distribution matching approach for collision avoidance with previous non-parametric and Gaussian approximated methods of reactive collision avoidance.

I. Introduction

Collision avoidance under uncertainty has been well studied in existing literature [1], [2], [3]. Most of them assume Gaussian perturbation in estimating robot and obstacles' state and motion commands executed by the robot for collision avoidance. Recent works like [4], [5] are capable of planning and control under non-Gaussian noise motion and perception noise. Some recent approaches like [6], [7], [2] specifically deal with reactive collision avoidance in a similar vein to our current work. These methods highlight the reduction in collision probability and control effort achieved by adopting a more sophisticated vies of the underlying uncertainty. However, they do not provide a fine-grained analysis of how distributional characteristics like bias affect collision avoidance and how we can leverage it to reduce collision probabilities. The bias we refer to can be described through figure 1

1. Robotics Research Center, IIIT Hyderabad, India 2. University of Tartu, Estonia. This work was supported in part by the European Social Fund via ICT programme measure, Estonian Center of Excellence in IT (EXCITE) funded by the European Regional Development Fund, grant PSG753 from Estonian Research Council, and Artificial Intelligence & Robotics Estonia (AIRE), the Estonian candidate for European Digital Innovation Hub, funded by the Ministry of Economic Affairs and Communications in Estonia.

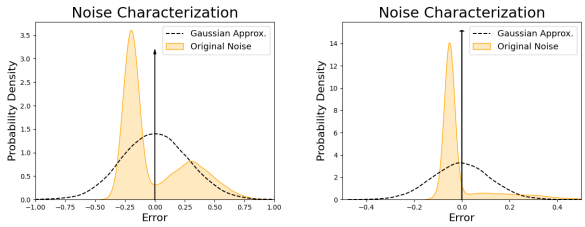

Fig. 1: Examples of Non-Gaussian distribution and their Gaussian approximations.

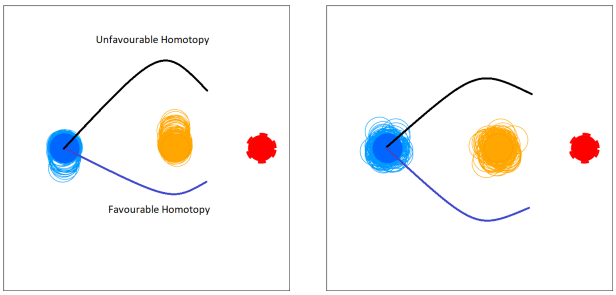

Fig. 2: Left and right figures show collision avoidance under Non-Gaussian and Gaussian approximation of uncertainty respectively

wherein a bi-modal distribution obtained from a commodity GPS is not mean-centered. In other words, the presence of the second mode provides for an unequal spread on either side of the mean, unlike its Gaussian approximation. As shown in Fig.2, the unequal spread of distribution mass naturally creates a notion of favorable and unfavorable homotopies. The former corresponds to the side where there would be less overlap between robot and obstacles uncertainty. In this paper, we fill this current knowledge gap by analyzing in diverse ways why a particular control action is chosen for a given obstacle configuration and how they depend on the nature of the underlying uncertainty and any approximation we make on it.

II. Related Work

Chance constrained optimization has emerged as the popular paradigm and framework for collision avoidance under uncertainty [1], [2], [8]. Most such formulations model the original distribution to be a Gaussian and resort to linearization [1] or manage closed form solutions when the collision avoidance constraints can be posed as convex or affine constraints [9], [10]. Recently there has been a growing trend towards non-parametric chance constraints [6], [7], [2] acknowledging that most sensor noise are more often non-parametric [6], [5]. These methods have not indulged in an analysis that pinpoints how and why non parametric model-

ing actually is beneficial. We fill this gap in this paper by providing detailed empirical analysis as to how the inherent bias prevalent in non-parametric noise is leveraged which does not happen in case Gaussian approximation. We are also contributing to our prior works [6], [11] by circumventing the need for estimating the desired distribution. We also extend [12] to reactive navigation in dynamic environments.

III. Preliminaries and Problem Formulation

Details regarding symbols and notations can be found at section III of [13]. Descriptions of the common symbols used in this paper can be found in Table I of [13].

A. Motion and Perception Model

We assume that the robot has the following discrete time stochastic motion model, wherein Δt represents the time duration between two consecutive steps.

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \mathbf{v}_t \Delta t, \theta_{t+1} = \theta_t + \omega_t \Delta t, \tag{1}$$

$$\mathbf{v}_t = \begin{bmatrix} v_t \cos(\theta_t + \omega_t \Delta t) \\ v_t \sin(\theta_t + \omega_t \Delta t) \end{bmatrix}, \tag{2}$$

$$\begin{bmatrix} v_t \\ \omega_t \end{bmatrix} = \overbrace{\begin{bmatrix} \overline{v}_t \\ \overline{\omega}_t \end{bmatrix}}^{\mathbf{u}_t} + \boldsymbol{\epsilon} \tag{3}$$

$$\mathbf{x}_{o,t+1} = \mathbf{x}_{o,t} + \mathbf{v}_{o,t} \Delta t \tag{4}$$

B. Velocity Obstacle Constraints

In deterministic noise-less setting, reactive collision avoidance between disk shaped robots and obstacles is often formulated in terms of velocity obstacle (VO) [14] constraints defined in the following manner:

$$f(\cdot) \leq 0 : \frac{(\mathbf{r}^T \mathbf{v})^2}{\|\mathbf{v}\|^2} - \|\mathbf{r}\|^2 + R^2 \leq 0, \forall j \tag{5a}$$

$$\mathbf{r} = \mathbf{x}_t - \mathbf{x}_{o,t}, \quad \mathbf{v} = \mathbf{v}_t - \mathbf{v}_{o,t} \tag{5b}$$

where, R represents the combined radii of the robot and the obstacle. For the ease of exposition in the latter sections, we formulate the VO constraints for a single obstacle. Extension to multiple obstacles is trivial.

C. Reactive Navigation Through CCO

We formulate one-step reactive navigation in uncertain environment as the following CCO:

$$\min_{\mathbf{u}_t} w_1 \|\overline{\mathbf{v}}_t - \mathbf{v}_d\|_2^2 + w_2 \|\overline{\mathbf{u}}_t\|^2 \tag{6a}$$

$$\Pr(f(\mathbf{x}_t, \theta_t, \mathbf{u}_t, \mathbf{x}_{o,t}, \mathbf{v}_{o,t}) \leq 0) \geq \eta, \ \forall j, \ \mathbf{u}_t \in \mathcal{C} \tag{6b}$$

$$\overline{\mathbf{v}}_t = \begin{bmatrix} \overline{v}_t \cos(\overline{\theta}_t + \overline{\omega}_t \Delta t) \\ \overline{v}_t \sin(\overline{\theta}_t + \overline{\omega}_t \Delta t) \end{bmatrix}, \tag{7}$$

The first term in the cost function (6a) ensures that the nominal velocity is aligned with some desired velocity vector $\mathbf{v}_d$ [15]. The $\mathcal{C}$ represents the set of feasible control inputs and we assume that it is convex formed by affine constraints on $\overline{v}_t, \overline{\omega}_t$. The set of inequalities (6b) represent the so-called chance constraints [7].

The main computational challenge in solving (6a)-(6b) stems from the computational intractability of chance constraints. Thus, existing works heavily focus on replacing (6b) with other feasible options. However, most of the existing reformulations assume that the underlying uncertainty is Gaussian [16]. Since our aim in this paper is to analyze the effect of Gaussian approximation, we next present our reactive planner that can work with arbitrary uncertainty distribution.

IV. Methods

At an intuitive level CCO (6a)-(6b) has the following interpretation [7]. We seek to compute a nominal control $\overline{\mathbf{u}}_t$ that modifies the shape of the distribution of $f(.)$ in a way that most of its mass lies on the left of the line $f_j(.) = 0$. An alternate interpretation can be derived by defining a function h in the following manner.

$$h(\mathbf{x}_t, \theta_t, \mathbf{u}_t, \mathbf{x}_{o,t}, \mathbf{v}_{o,t}) = \max(0, f(\mathbf{x}_t, \theta_t, \mathbf{u}_t, \mathbf{x}_{o,t}, \mathbf{v}_{o,t})) \tag{8}$$

As clear, $h(.)$ measures constraint violation. It is zero if the VO constraints are satisfied and equal to $f(.)$ otherwise. In the stochastic setting where $\mathbf{x}_t, \mathbf{u}_t, \mathbf{x}_{oj,t}, \mathbf{v}_{oj,t}$ are random variables, $h(.)$ defines the distribution of constraint violations.

With respect to (8), we can interpret CCO as the problem of finding an appropriate control input $\overline{\mathbf{u}}_t$ such that the distribution of $h(.)$ becomes similar to that of a Dirac-Delta. Using this interpretation, we can reformulate (6a)-(6b) in the following manner:

$$\min_{\overline{\mathbf{u}}_t} l_{dist}(p_h, p_\delta) + w_1 \|\overline{\mathbf{v}}_t - \mathbf{v}_d\|_2^2 + w_2 \|\overline{\mathbf{u}}_t\|^2 \tag{9}$$

$$\mathbf{u}_t \in \mathcal{C}, \tag{10}$$

where p_h, p_δ represents the probability distribution of $h(.)$ and Dirac-Delta respectively. The function l_{dist} measures the similarity between p_h, p_δ and it decreases as the distribution becomes similar. One possible option for l_{dist} is the KL divergence. However, it cannot operate at purely sample level and requires the parametric form of the distributions to be known. Thus, we define l_{dist} as MMD between p_h and p_δ defined in the following manner.

$$l_{dist}(\mathbf{u}_t) = \overbrace{\|\mu_{p_h}(\mathbf{u}_t) - \mu_{p_\delta}\|_2^2}^{MMD}, \tag{11}$$

where, μ_{p_h} and μ_{p_δ} represent the RKHS embedding of p_h and p_δ respectively.

We solve (9)-10 through a simple control sampling approach. We draw several samples of $\mathbf{u}_t$ from a uniform distribution and then evaluate the cost (9) on them. Subsequently, we choose the sample corresponding to the lowest cost. Our control sampling relies on efficient evaluation of MMD term to retain online performance. The MMD value for a given $\mathbf{u}_t$ can be reduced to computing matrix-matrix products. This is done in section IVB of [13].

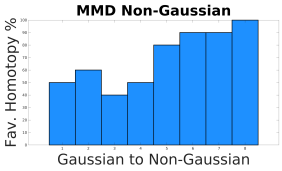

(a) This bar plot depicts the frequency of choosing favorable homotopy with MMD Non-Gaussian in a single obstacle benchmark shown in Fig.4b,4c

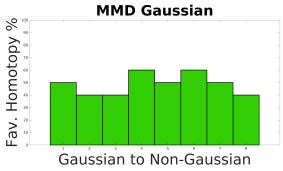

(b) This bar plot depicts the frequency of choosing favourable homotopy for MMD Gaussian in case of single obstacle benchmark shown in Fig.4b,4c

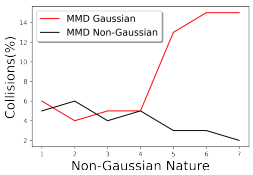

(c) Effect of non-gaussian nature collision probability

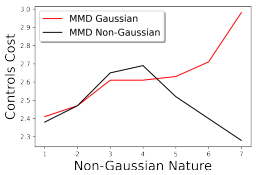

(d) Effect of non-gaussian nature on control costs.

Fig. 3: Quantitative Analysis on non-gaussian nature of distribution. The x-axis shows the distribution number from Fig. 5 from [13]

V. Results and Discussion

The implementation details and baselines are explained at the beginning of section V of [13]. The validation of our approach can be seen at section VA from [13]. Further results can be found at https://github.com/anishgupta31296/MMD-with-Dirac-Delta-Distribution.

A. Analyzing the Choice of Homotopies

We consider a benchmark with a single obstacle as shown in Fig.4 to analyze two key questions. First, how is the choice of homotopy related to the distribution of constraint violations for a biased non-Gaussian distribution and its Gaussian approximation. Second, we intend to study the effectiveness of our MMD Non-Gaussian approach in ensuring the selection of favorable homotopies during collision avoidance. To these ends, we sampled two control actions for the scenario shown in Fig.4 which results in the robot passing the obstacle from different sides. Clearly, Fig.4a is the favorable homotopy in this scenario while that shown in Fig.4d leads to a large overlap between the robot and obstacle position uncertainty. Fig.4b shows the distribution of constraint violations for the control input that leads to the favorable homotopy for the true non-Gaussian distribution. It can be seen that the distribution of violation is very close to the ideal Dirac-Delta distribution. Now, contrast this with Fig.4e that recreates the constraint violation distribution for the control input leading to unfavorable homotopy. We can clearly see a stark difference between Fig.4b and 4e. The reason for this can be seen in Fig4c and 4f which presents the distribution of constraint violations under Gaussian approximation of the noise. As it can be seen, both favorable and un-favorable homotopy shows similar spread of the distribution mass to the right of zero. As a result, it is not possible to reliably distinguish between favorable and unfavorable homotopies. This leads us to hypothesize that any planner that can capture the true distribution of constraint violation for a given control input can easily distinguish between a favorable and unfavorable homotopy.

To further strengthen our claims, we design one more experiment. In Fig. 5 of [13], we take a Gaussian distribution and then gradually make it more and more biased and multi-modal. We simulate the single obstacle avoidance benchmark of Fig.4 for all these noise distributions added to motion and perception. We perform 100 Monte-Carlo runs for each noise distribution using our MMD Non-Gaussian planner. Fig.3a shows the percentage of times the robot chooses homotopy of Fig.4a over that of Fig.4d. When the actual noise is Gaussian, the robot randomly chooses either homotopy. But as the noise becomes more and more non-Gaussian, we can clearly see a pattern emerge where the favorable homotopy is overly preferred by our planner. In contrast when we make a Gaussian approximation of the true uncertainty, this pattern is lost, as shown in Fig.3b. Under Gaussian approximation, the robot always chooses the homotopies randomly.

Fig.3c and 3d co-relates the right choice of homotopy to collision percentages and control cost. When the underlying noise is Gaussian, both MMD Non-Gaussian and MMD Gaussian performs similar. But as the distribution departs from Gaussian assumptions, the former outperforms the latter in both collision-rate and control costs.

B. Quantitative Comparisons

In this section, we compare our MMD Non-Gaussian formulation with MMD Gaussian, KLD and PVO baselines defined at the beginning of section V of [13]. The comparisons are shown in the bar plots of Figure 6. Fig.5 presents the trajectories observed in a 5 obstacle benchmark for all the approaches. Our MMD Non-Gaussian is able to leverage the bias of the distribution and guide the robot towards homotopies that goes between the obstacles but yet has minimal overlap of robot and obstacle position uncertainty. In contrast, both MMD Gaussian and PVO that works with Gaussian approximation of noise forces the robot to take a larger detour. This is because the Gaussian approximation over-approximates the spread of the uncertainty on either side of the robot mean position. The KLD method shows a very similar approach since it fits a complicated a GMM to the motion and perception noise.

Figure 6a compares over L_2 norm of control change over two consecutive instances $\|\mathbf{u}_t - \mathbf{u}_{t-1}\|_2^2$ which can be used to infer the smoothness of a collision avoidance maneuver. Our approach has the lowest change while all other baselines have similar trends. Fig.6b shows the comparison between the deviation that the robot exhibits from an optimal straight line path to the goal. On an average our approach is 72.84% better than all the other baselines. Finally, we compare the collision probabilities in Fig.6c. Our approach maintains a value of 5% or less, compared to baselines' $11 - 28\%$.

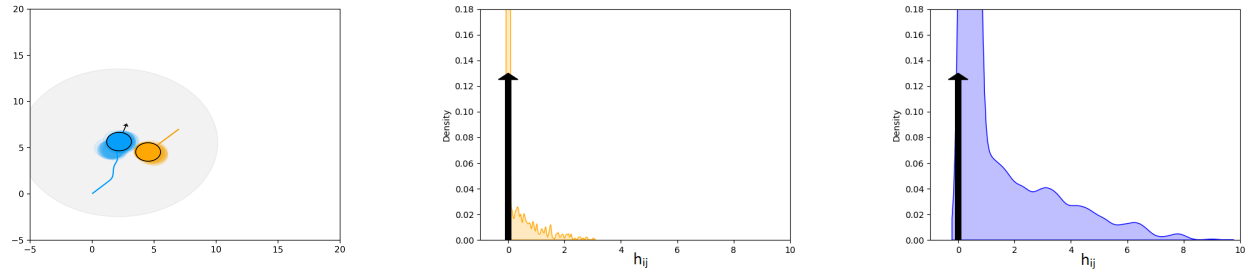

(a) Favorable homotopy chosen by the robot by leveraging distribution bias (b) Corresponding VO constraint violation distribution (c) Corresponding VO constraint violation distribution under Gaussian approximation

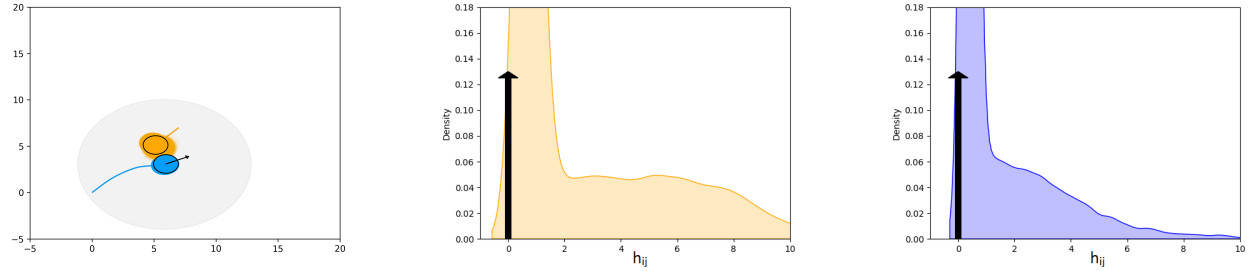

(d) Unfavorable homotopy chosen by the robot. (e) Corresponding VO constraint violation distribution (f) Corresponding VO constraint violation distribution under Gaussian approximation

Fig. 4: These figures illustrate how better favourable homotopy selection will lead to better distribution matching and hence, larger number of samples will be avoided

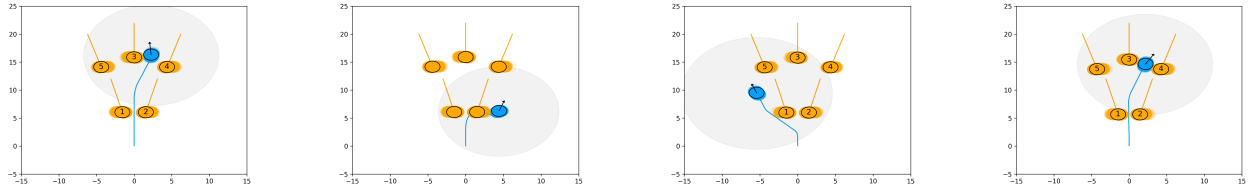

(a) MMD Non-Gaussian (b) MMD Gaussian (c) PVO (d) KLD

Fig. 5: Collision avoidance using MMD Non-Gaussian and various baselines for 5 obstacle case

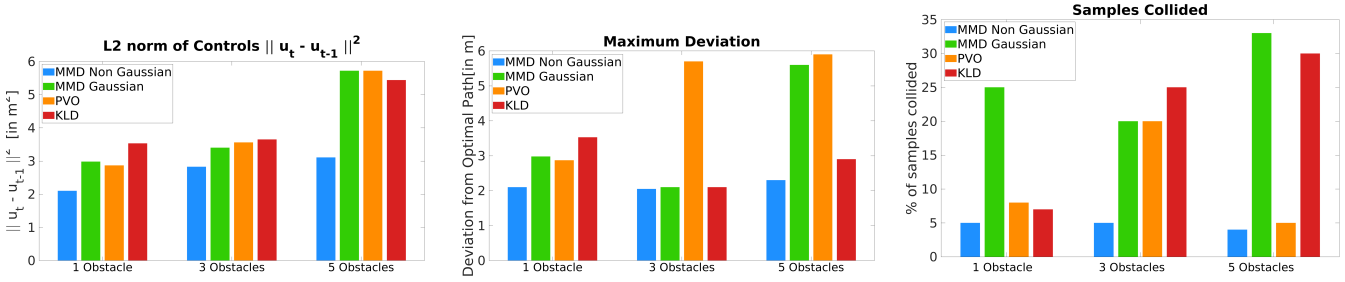

(a) Control Costs comparison (b) Deviation from optimal path comparison (c) Number of colliding samples

Fig. 6: Quantitative comparison with baselines. Our approach MMD Non-Gaussian outperforms other approaches in smoothness (a), deviation from straight-line path (b) and collision probability (c) metric.

REFERENCES

[1] H. Zhu and J. Alonso-Mora, "Chance-constrained collision avoidance for mavs in dynamic environments," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 776–783, 2019.

[2] J. S. Park and D. Manocha, "Efficient probabilistic collision detection for non-gaussian noise distributions," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1024–1031, 2020.

[3] B. Kluge and E. Prassler, "Recursive agent modeling with probabilistic velocity obstacles for mobile robot navigation among humans," in *Autonomous Navigation in Dynamic Environments*. Springer, 2007, pp. 121–134.

[4] D. M. Rosen and J. J. Leonard, "Nonparametric density estimation for learning noise distributions in mobile robotics," in *1st Workshop on Robust and Multimodal Inference in Factor Graphs, ICRA*, 2013.

[5] A. Majumdar and V. Pacelli, "Fundamental performance limits for sensor-based robot control and policy learning," *arXiv preprint arXiv:2202.00129*, 2022.

[6] S. N. J. Poonganam, B. Gopalakrishnan, V. S. S. B. K. Avula, A. K. Singh, K. M. Krishna, and D. Manocha, "Reactive navigation under non-parametric uncertainty through hilbert space embedding of probabilistic velocity obstacles," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2690–2697, 2020.

[7] B. Gopalakrishnan, A. K. Singh, K. M. Krishna, and D. Manocha, "Solving chance-constrained optimization under nonparametric uncertainty through hilbert space embedding," *IEEE Transactions on Control Systems Technology*, 2021.

[8] B. Gopalakrishnan, A. K. Singh, and K. Krishna, "Closed form characterization of collision free velocities and confidence bounds for non-holonomic robots in uncertain dynamic environments," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 4961–4968.

[9] S. Boyd and A. Mutapcic, "Stochastic subgradient methods," *Lecture Notes for EE364b, Stanford University*, 2008.

[10] L. Blackmore, M. Ono, A. Bektassov, and B. C. Williams, "A probabilistic particle-control approximation of chance-constrained stochastic predictive control," *IEEE transactions on Robotics*, vol. 26, no. 3, pp. 502–517, 2010.

[11] U. K. R. Nair, A. Gupta, D. A. S. Kiran, A. Shrihari, V. Shah, A. K. Singh, and K. M. Krishna, "Non holonomic collision avoidance under non-parametric uncertainty: A hilbert space approach," in *2021 European Control Conference (ECC)*, 2021, pp. 675–681.

[12] S. S. Harithas, R. D. Yadav, D. Singh, A. K. Singh, and K. M. Krishna, "Cco-voxel: Chance constrained optimization over uncertain voxel-grid representation for safe trajectory planning," *arXiv preprint arXiv:2110.02904*, 2021.

[13] A. Gupta, A. K. Singh, and K. M. Krishna, "Leveraging distributional bias for reactive collision avoidance under uncertainty: A kernel embedding approach," *arXiv preprint arXiv:2208.03038*, 2022.

[14] P. Fiorini and Z. Shiller, "Motion planning in dynamic environments using velocity obstacles," *The International Journal of Robotics Research*, vol. 17, no. 7, pp. 760–772, 1998.

[15] J. van den Berg, Ming Lin, and D. Manocha, "Reciprocal velocity obstacles for real-time multi-agent navigation," in *2008 IEEE International Conference on Robotics and Automation*, May 2008, pp. 1928–1935.

[16] H. Zhu and J. Alonso-Mora, "Chance-constrained collision avoidance for mavs in dynamic environments," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 776–783, 2019.