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

Thesis submitted in partial fulfillment
of the requirements for the degree of

Master of Science in
Computer Science and Engineering
by Research

by

ANISH GUPTA
2020701025
anish.gupta@research.iiit.ac.in

International Institute of Information Technology
Hyderabad - 500 032, INDIA
January 2024

International Institute of Information Technology
Hyderabad, India

CERTIFICATE

It is certified that the work contained in this thesis, titled "Leveraging Distributional Bias For Reactive Collision Avoidance under Uncertainty: A Kernel Embedding Approach" by ANISH GUPTA, has been carried out under my supervision and is not submitted elsewhere for a degree.

Date	Adviser: Prof. K MADHAVA KRISHNA

To My Family and Teachers

Acknowledgments

I would like to express my deepest gratitude to my guide, Professor K Madhava Krishna, whose unwavering guidance, support, and mentorship have been instrumental in shaping this thesis. His expertise, insightful feedback, and encouragement throughout this research journey have been invaluable, and I am immensely grateful for his commitment to my academic and personal growth.

I am also indebted to my Professor Arun Kumar Singh from University of Tartu, Estonia, for their significant contributions and collaboration on this thesis. Their international expertise and unique perspective have added a valuable dimension to my research, enriching the overall quality of this work. I am grateful for the opportunity to learn from them and for their unwavering support, even across borders and time zones.

Furthermore, I extend my appreciation to the entire faculty of Robotics Research Center for fostering an intellectually stimulating environment that has nurtured my academic curiosity and passion for research.

I am indebted to my family and friends for their unyielding encouragement, love, and belief in my capabilities, which have been a constant source of motivation during challenging times.

Finally, I would like to express my heartfelt thanks to all the individuals who participated in this research, as well as to the countless anonymous reviewers whose valuable feedback has contributed to the refinement of this thesis.

It is with profound gratitude that I acknowledge the support and contributions of all those who have played a role in shaping this thesis and my academic journey. Your influence has been immeasurable, and I am humbled by the opportunity to share this work with the world.

Thank you.

Abstract

In the field of robotics, the accurate modeling of uncertainty in the robot's motion and perception is crucial for effective collision avoidance. Many commodity sensors exhibit non-Gaussian noise characteristics, yet existing approaches often assume Gaussian uncertainty to ensure computational tractability. This thesis addresses the gap between non-Gaussian uncertainty and collision avoidance by developing a framework that leverages the distributional characteristics of motion and perception noise.

We propose a novel approach that interprets reactive collision avoidance as a distribution matching problem between collision constraint violations and the Dirac Delta distribution. To ensure fast reactivity, we embed each distribution in a Reproducing Kernel Hilbert Space and reformulate the distribution matching as the minimization of the Maximum Mean Discrepancy (MMD). By exploiting the insight that evaluating the MMD reduces to matrix-matrix products, we develop a simple control sampling approach for reactive collision avoidance with dynamic and uncertain obstacles.

Furthermore, this thesis advances the state-of-the-art in two key aspects. Firstly, we conduct an extensive empirical study to demonstrate that our planner can effectively infer distributional bias from sample-level information. This insight enables the planner to guide the robot towards good homotopy, utilizing the distributional characteristics of motion and perception noise. In contrast, we highlight how a Gaussian approximation of uncertainty can lead to loss of bias estimation and guide the robot towards unfavorable states with high collision probabilities. Secondly, we compare our proposed distribution matching approach with previous non-parametric and Gaussian approximated methods of reactive collision avoidance. Through tangible comparative advantages, we showcase the superior performance of the distribution matching approach.

In summary, this thesis presents a comprehensive framework that addresses the challenge of non-Gaussian uncertainty in collision avoidance. By leveraging the distributional characteristics of motion and perception noise, our approach provides a more accurate and effective method for reactive collision avoidance with dynamic and uncertain obstacles. The empirical study and comparative evaluations demonstrate the advantages of the proposed distribution matching approach over previous methods. This research contributes to advancing the understanding and applicability of collision avoidance strategies, particularly in the context of non-holonomic motion and non-Gaussian uncertainty.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Robotics, as a multidisciplinary field, has witnessed significant advancements in recent years. From industrial automation to healthcare assistance, robots have become increasingly prevalent in various domains, revolutionizing the way tasks are performed and challenging traditional notions of human-machine interaction. The quest for intelligent, autonomous machines capable of perceiving, reasoning, and acting in complex environments has fueled rapid progress in robotics research.

One key area of focus within robotics is the development of autonomous systems. These systems aim to imbue robots with the ability to operate independently, making decisions and executing tasks without continuous human intervention. Autonomy enables robots to adapt to dynamic environments, handle uncertainties, and accomplish complex missions with precision and efficiency.

Within the realm of autonomous systems, a significant application area is autonomous vehicles. The emergence of self-driving cars, unmanned aerial vehicles, and autonomous underwater vehicles has garnered considerable attention and investment. The promise of autonomous vehicles lies in their potential to transform transportation systems, offering increased safety, improved mobility, and reduced environmental impact.

However, the realization of autonomous vehicles is not without challenges. One of the foremost concerns is ensuring collision avoidance, a critical aspect for safe and reliable operation. In dynamic and unpredictable environments, autonomous vehicles must be capable of detecting and responding to obstacles, pedestrians, and other vehicles in real-time. Effective collision avoidance strategies are essential to mitigate the risk of accidents and enable seamless integration of autonomous vehicles into existing transportation networks.

In the pursuit of robust collision avoidance, researchers have explored various techniques and algorithms. Traditionally, many approaches have relied on assumptions of Gaussian noise and linear models, providing tractability but potentially oversimplifying the true complexity of real-world scenarios. Acknowledging the non-Gaussian noise characteristics exhibited by many commodity sensors, recent research has sought to leverage the distributional properties of motion and perception noise to improve collision avoidance performance.

By considering the biases and unequal spreads in uncertainty distributions, novel approaches have emerged that can guide robots towards favorable states while avoiding collisions. These advancements not only enhance the safety and efficiency of autonomous systems but also have broader implications for the field of robotics as a whole.

In this thesis, we contribute to the ongoing research in collision avoidance by investigating the impact of non-Gaussian uncertainty and distributional characteristics on the performance of robotic systems. Our work aims to improve the understanding of how biases and non-Gaussian noise affect collision avoidance outcomes in diverse robotics applications. We develop novel algorithms that leverage distribution matching and Reproducing Kernel Hilbert Space to enable efficient reactive collision avoidance with dynamic and uncertain obstacles.

Through empirical studies and comparative analyses, we demonstrate the advantages of our proposed methods over traditional approaches that rely on Gaussian approximations. Furthermore, we highlight the broader significance of our findings in the context of autonomous vehicles, where collision avoidance is of paramount importance for safe and reliable operation.

By advancing the knowledge and techniques in collision avoidance with non-Gaussian uncertainty in robotics, this thesis contributes to the broader field of autonomous systems, paving the way for more capable and intelligent robots in various domains.

1.1 Collision Avoidance Under Uncertainty

Collision avoidance in robotics is a critical task that involves ensuring the safe navigation of robots in dynamic environments, avoiding collisions with obstacles and other agents. However, real-world scenarios often introduce uncertainties in the robot's motion and perception, making collision avoidance a challenging problem. The literature on collision avoidance under uncertainty has been extensive, with many studies assuming Gaussian perturbations for estimating the robot and obstacles' states, as well as the robot's motion commands [1–3]. The Gaussian approximation is commonly employed due to its computational tractability and the resulting efficient convex structures in the problem [4].

While Gaussian approximations are suitable for many cases, they may not accurately capture the complexities of uncertainty when the underlying noise deviates significantly from Gaussian distributions. In such instances, the reliance on Gaussian assumptions can lead to conservative

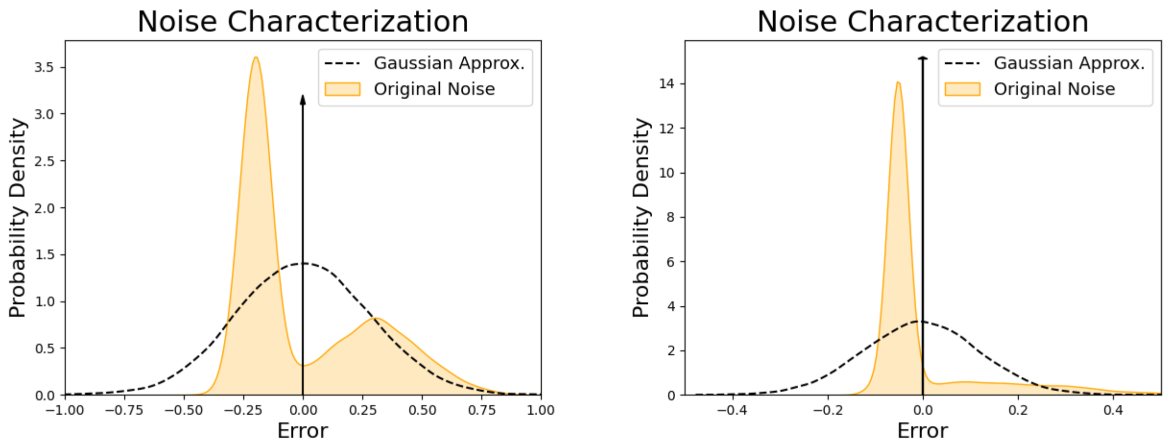

Figure 1.1: Examples of Non-Gaussian distribution and their Gaussian approximations. The majority of the mass of the true distribution is shifted with respect to the mean of the Gaussian approximation. We refer to it as the distribution bias.

estimates of the feasible space, adversely affecting the efficiency of collision avoidance planning.

Recent advancements have been made to address the limitations of Gaussian-based approaches by considering non-Gaussian noise models for motion and perception [5, 6]. These works enable planning and control under more realistic non-Gaussian noise, offering the potential for more accurate collision avoidance strategies. Additionally, specific attention has been given to reactive collision avoidance [2, 7, 8], where strategies aim to avoid collisions in real-time while maintaining efficiency and control effort. These approaches highlight the benefits of adopting more sophisticated views of underlying uncertainty, leading to reduced collision probabilities and optimized control actions.

However, despite these notable contributions, existing approaches have not provided a fine-grained analysis of how distributional characteristics, such as bias, impact collision avoidance and how this knowledge can be leveraged to further reduce collision probabilities. Bias, as exemplified in 1.1 with a bi-modal distribution from a commodity GPS, can result in unequal spread on either side of the mean, deviating from the Gaussian approximation. This unequal spread inherently creates a notion of favorable and unfavorable homotopies, as depicted in 1.2, where regions with less overlap between robot and obstacle uncertainties are perceived as favorable homotopies. To the best of our knowledge, current methods have not explicitly ensured that robots consistently choose favorable homotopies with high probability while avoiding obstacles.

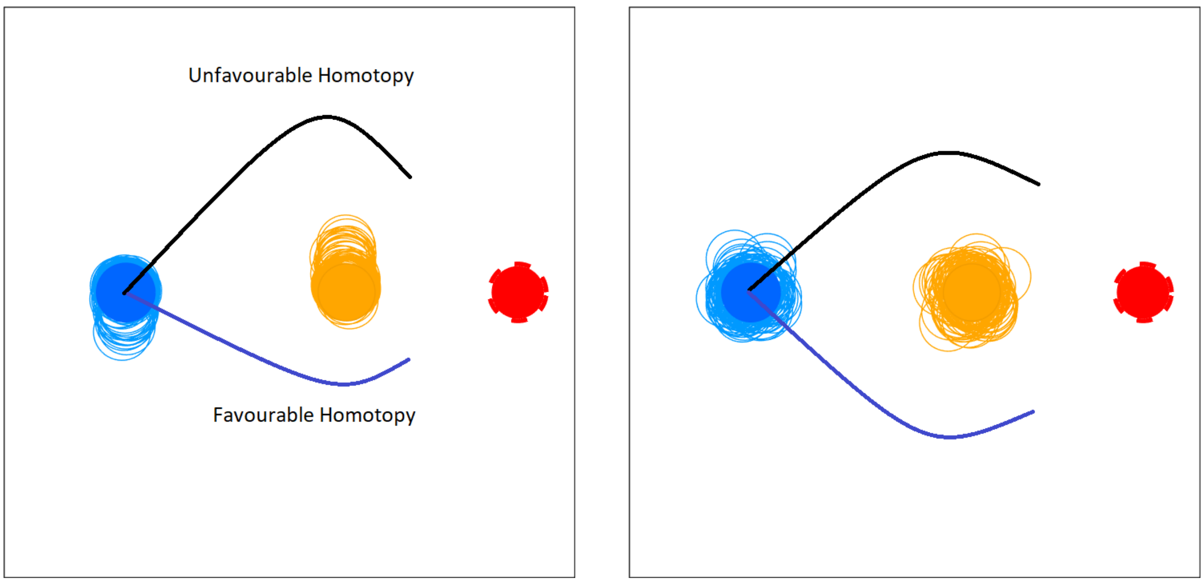

Figure 1.2: Left figure shows collision avoidance under non-Gaussian motion and perception noise. The goal position is shown in red. The robot (blue) can choose to avoid the obstacle (orange) from either left or right. However, due to the presence of bias in the motion and perception noise, one of the homotopies shown in blue becomes more favorable. Our objective in this paper is to develop reactive planners than can guide the robots towards favorable homotopies. The right figure presents the situation under Gaussian approximation of the noise. In this case, either homotopy erroneously appear equally good (or worse). Thus, it is quite likely that Gaussian approximation will lead the robot unfavorable positions with high collision probability

Table 1.1

Symbol	Description
$(\mathbf{x}_t, \mathbf{v}_t)$	Position and linear velocity of the robot at time t
(θ_t, ω_t)	Heading and angular velocity of the robot at time t
$(\mathbf{x}_{o,t}, \mathbf{v}_{o,t})$	Position and velocity of the obstacle at time t
$\mathbf{u}_t$	Control input to the robot at time t
$f(\cdot) \leq 0$	VO constraints for j^{th} obstacle
p_f	Distribution of $f(.)$ under motion and perception uncertainty
η	Probability of collision avoidance

In this thesis, we aim to bridge this knowledge gap and enhance the understanding of collision avoidance under non-Gaussian uncertainty. Our approach involves analyzing in diverse ways why specific control actions are chosen for given obstacle configurations and how these decisions are influenced by the nature of the underlying uncertainty and any approximations made. By considering the distributional characteristics, particularly bias, we seek to develop novel algorithms that can effectively infer the favorable homotopies, guiding robots to make safer and more efficient collision avoidance decisions in real-world environments.

Through extensive empirical studies and comparative evaluations, we demonstrate the efficacy of our proposed approach, providing valuable insights into collision avoidance under uncertainty. Our research contributes to the advancement of robotics and autonomous systems, fostering safer and more reliable robot navigation in complex and uncertain environments. By addressing the challenges posed by non-Gaussian uncertainty, we aim to pave the way for the seamless integration of robots into real-world applications and propel the field of collision avoidance to new heights of performance and innovation.

1.2 Preliminaries

Symbols and Notations: We represent scalars as normal case small font letters and use the bold font variant for vectors. We use bold-font upper case letters to denote matrices. We use subscript t to denote the time-stamp of a variable. The notation $\|\cdot\|$ denotes the Euclidean norm of vector/matrices. We use $\overline{c}$ to denote some nominal/noise-free value of a random variable c. The symbol $\Pr(\cdot)$ denotes the probability of an event, while $p(.)$ represents the probability distribution function. Some of the commonly used symbols and notations are summarized in the table 1.1 while some are also defined in their first place of use.

1.2.1 Motion Model

In our investigation, we consider a discrete-time stochastic motion model for the robot, where Δt denotes the time interval between consecutive steps.

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \mathbf{v}_t \Delta t, \theta_{t+1} = \theta_t + \omega_t \Delta t, \tag{1.1}$$

$$\mathbf{v}_t = \begin{bmatrix} v_t \cos(\theta_t + \omega_t \Delta t) \\ v_t \sin(\theta_t + \omega_t \Delta t) \end{bmatrix}, \tag{1.2}$$

$$\begin{bmatrix} v_t \\ \omega_t \end{bmatrix} = \overbrace{\begin{bmatrix} \overline{v}_t \\ \overline{\omega}_t \end{bmatrix}}^{\mathbf{u}_t} + \boldsymbol{\epsilon} \tag{1.3}$$

In the stochastic setting, $\mathbf{x}_t, \theta_t, v_t, \omega_t$ are all random variables with unknown probability distribution. To simplify the technical exposition, we assume that these variables have a nominal noise-free value corrupted by an additive disturbance. For example, as shown in (1.3), the control consists of deterministic command of linear ($\overline{v}_t$) and angular ($\overline{\omega}_t$) velocity corrupted by $\boldsymbol{\epsilon}$. Although, the probability distribution of the disturbance is not known, we assume to have access to the samples drawn from it. We also assume that a Particle filter like set-up is in place that bounds the uncertainty in position at each time step.

We represent obstacles' motions through the following piece-wise straight line trajectory. Similar to robot motion model, we treat $\mathbf{x}_{oj,t}, \mathbf{v}_{oj,t}$ as random variables and we have access to only small number of sample realizations of these random variables.

$$\mathbf{x}_{o,t+1} = \mathbf{x}_{o,t} + \mathbf{v}_{o,t} \Delta t \tag{1.4}$$

1.2.2 Velocity Obstacle

In a deterministic noise-less scenario, reactive collision avoidance between disk-shaped robots and obstacles is frequently expressed using velocity obstacle(VO) [9] constraints, defined as follows:

$$f(\cdot) \leq 0 : \frac{(\mathbf{r}^T\mathbf{v})^2}{\mathbf{v}^2} - \mathbf{r}^2 + R^2 \leq 0, \forall j \tag{1.5a}$$

$$\mathbf{r} = \mathbf{x}_t - \mathbf{x}_{o,t}, \quad \mathbf{v} = \mathbf{v}_t - \mathbf{v}_{o,t} \tag{1.5b}$$

Here, R represents the combined radii of the robot and the obstacle. For clarity in subsequent sections, we formulate the VO constraints for a single obstacle, but extending to multiple obstacles is straightforward. However, in the stochastic setting, where the positions and velocities of the robot and obstacle are all random variables, the function f(.) effectively characterizes a distribution. As a result, we must appropriately modify the application of VO constraints to accommodate this situation, which will be elaborated on in the following discussion.

1.3 Navigating Uncertainties: Advancements in Non-Gaussian Collision Avoidance

1.3.1 Motivation

Handling non-parametric noise in robotics and autonomous systems is of paramount importance due to the inherent complexities and uncertainties present in real-world environments. In many practical scenarios, the motion and perception noise encountered by robots and sensors do not conform to simple Gaussian models. Assuming Gaussian noise distributions can lead to suboptimal decision-making, as it may overlook critical information contained within the true non-Gaussian noise characteristics. Embracing non-parametric noise models allows us to capture more realistic and intricate behaviors, enabling robots to make informed and adaptive decisions in dynamic and unpredictable surroundings as seen in Fig 1.3. By acknowledging the non-Gaussian nature of noise, we can design algorithms that effectively handle uncertainties, leading to more robust and reliable autonomous systems capable of navigating complex environments with higher levels of precision and safety.

Moreover, the consideration of non-parametric noise contributes to the identification of favorable homotopies in collision avoidance scenarios. Non-Gaussian noise distributions often exhibit unequal spreads and biases, resulting in non-symmetric distributions. As a consequence, certain regions in the environment may have less overlap between robot and obstacle uncertainties, creating favorable homotopies that allow safer and more efficient navigation. By leveraging distributional characteristics, such as bias, in the collision avoidance process, robots can intelligently select trajectories that increase the probability of avoiding collisions while optimizing their path to the goal. This fine-grained analysis of the underlying uncertainty and

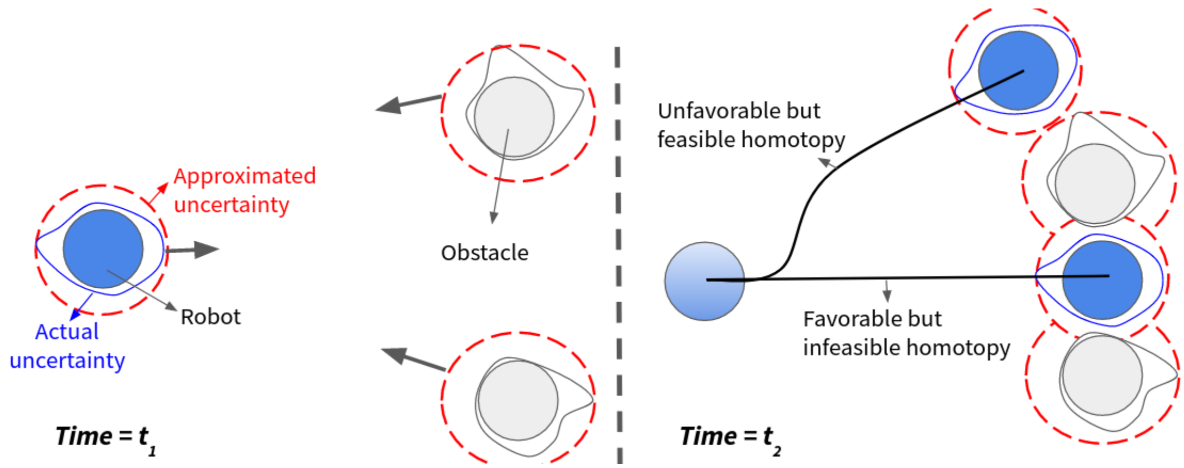

Figure 1.3: Examples showing the consequence of Gaussian approximation on homotopy selection. As we can see, on Gaussian approximation, the unfavourable homotopy is selected. This is because overapproximation of noise in a specific direction makes the favourable homotopy infeasible

its impact on decision-making allows robots to navigate around obstacles with higher success rates, ultimately leading to more favorable and efficient trajectories in complex and uncertain environments.

1.3.2 Related Works

Chance constrained optimization has emerged as a widely adopted paradigm and framework for addressing collision avoidance under conditions of uncertainty [1, 2]. While diverse variants of this problem exist, typical formulations seek to reformulate the initially intractable chance constraints into surrogate counterparts. In specific cases, such as demonstrated in [10], closed-form solutions are attainable. Often, these formulations model the original distribution as Gaussian and resort to linearization [1] or manage closed-form solutions when the collision avoidance constraints can be formulated as convex or affine constraints [11, 12]. Certain methods, like [13–15], have devised surrogates that provide tight approximations to chance constraints defined over non-linear inequalities (which may even be non-convex). Additionally, [3] proposes a Bayesian Decomposition framework for multi-robotic settings. However, all these algorithms share a fundamental assumption concerning the nature of uncertainty of the random variables involved (e.g., state and actuation of the robot), relying on Gaussian distributions for closed-form surrogates.

In recent times, there has been a growing inclination towards non-parametric chance constraints [2, 7, 8], recognizing that many sensor noises frequently exhibit non-parametric characteristics [6, 7]. These methods often present chance constraint optimization as a distribution matching problem, employing techniques such as the popular Kullback-Leibler (KL) distance or computing distances between distributions in their Hilbert Space embeddings. While showcasing promising outcomes in terms of various metrics, these methods have yet to delve into a thorough analysis that precisely identifies the reasons and implications of non-parametric modeling's benefits and the potential inadequacies associated with Gaussian approximations, particularly when it comes to control actions and resulting outcomes.

In this paper, we undertake the task of bridging this knowledge gap by conducting a detailed empirical analysis, exposing how inherent bias inherent in non-parametric distributions can pose challenges when approximated by parametric Gaussian noise. Moreover, we extend previous work [7] by circumventing the need to estimate the desired distribution in the distribution matching interpretation of chance-constrained optimization (CCO). Furthermore, our contributions include extending the scope of [16] to encompass reactive navigation in dynamic environments, thereby enhancing the applicability and practicality of our approach in real-world scenarios. Through comprehensive analysis and meticulous experimentation, we aim to shed light on the merits and limitations of non-parametric modeling, ultimately paving the way for more effective and reliable collision avoidance strategies in uncertain and dynamic robotic environments.

1.3.3 Contribution

Our research encompasses groundbreaking contributions rooted in the chance-constrained optimization (CCO) framework, encapsulating core advancements as follows.

Algorithmic Contribution: The fundamental premise of our work lies in harnessing bias in motion and perception noise to enhance planning efficiency. This necessitates the integration of reactive planners capable of operating amidst non-parametric uncertainty, inspired by our prior work [7], [8], [9], which reinterprets CCO as a distribution matching problem. Specifically, we reformulate CCO to identify suitable control inputs that minimize the discrepancy between the violation of velocity obstacle (VO) constraints and Dirac-Delta distributions. Leveraging distribution embedding in Reproducing Kernel Hilbert Space (RKHS), we devise the distribution matching cost using the Maximum Mean Discrepancy (MMD) measure. Additionally, leveraging the kernel trick, we streamline MMD evaluation to entail only a few matrix-matrix products, facilitating an uncomplicated control sampling approach for real-time reactive navigation.

Empirical Contribution: Our research elucidates the profound significance of retaining the true non-parametric nature of the distribution when computing motion plans. Importantly, we demonstrate that with an appropriate planner, such as the one proposed in this research, leveraging distributional bias substantially augments collision probability and control effort by deftly guiding robots towards favorable homotopies. By contrast, approximating uncertainty as Gaussian yields a planner's homotopy selection with an arbitrary nature, leading to heightened collision probabilities.

Benchmarking Contribution: A comprehensive benchmarking analysis of our planner against diverse baselines highlights substantial improvements in collision probabilities and control costs. The first baseline [10] adopts Gaussian approximation of motion and perception uncertainty, allowing for tractable reformulation of chance constraints. The second baseline follows the same distribution matching interpretation of CCO as ours but employs Gaussian Mixture Model for uncertainty fitting and utilizes Kullback Liebler Divergence (KLD) to construct distribution matching cost. Our final baseline constitutes an ablation study wherein our MMD-based approach is employed, but with Gaussian approximation of uncertainty.

In conclusion, our research not only introduces novel insights into non-Gaussian collision avoidance but also establishes a cutting-edge framework for integrating distributional characteristics into robotic motion planning. The empirical analysis and benchmarking comparisons underscore the superiority of our proposed approach, paving the way for more robust and efficient collision avoidance strategies in real-world autonomous systems. By embracing non-parametric noise models and leveraging distributional characteristics, our research pushes the boundaries of autonomous robotics, revolutionizing collision avoidance capabilities in complex and uncertain environments.

1.4 Organization of the Thesis

The thesis is organized into four chapters, each contributing to a comprehensive understanding of collision avoidance under non-parametric uncertainty and its benefits in robotic motion planning.

- **Chapter 1:** This introductory chapter serves as the foundation, starting with an introduction to the problem of collision avoidance and the significance of handling non-Gaussian noise. Preliminary concepts, including the motion model and velocity obstacle, are discussed to establish a solid basis for the subsequent chapters. The chapter delves into related works, highlighting existing approaches in collision avoidance and their limitations. The motivation behind the research and the specific contributions of the thesis are also presented, setting the stage for the subsequent investigations.

- **Chapter 2:** The chapter is dedicated to the problem formulation and the proposed novel approach for reactive collision avoidance. It elaborates on how chance-constrained optimization (CCO) is reformulated as a distribution matching problem to handle non-parametric noise effectively. The Maximum Mean Discrepancy (MMD) cost and the reduced sets method are detailed as crucial components of the proposed approach. A key innovation involves introducing Dirac Delta as the desired distribution to guide the robot towards favorable homotopies. This chapter lays the technical groundwork for the empirical evaluations presented later in the thesis.

- **Chapter 3:** This chapter focuses on the validation results and empirical analysis. The proposed method is put to the test through various simulations and experiments in dynamic environments. The chapter provides insights into homotopy selection and the impact of distributional characteristics on collision avoidance. Qualitative results are presented to demonstrate the efficacy of the proposed approach compared to traditional Gaussian-based methods and other non-parametric methods. The analysis showcases the superiority of the proposed approach in terms of collision probabilities and control efforts

- **Chapter 4:** This chapter concludes the thesis, summarizing the key findings and contributions. The implications of the empirical results and the significance of the novel Dirac Delta approach are discussed. The chapter also outlines potential future directions for research, suggesting ways to enhance and extend the proposed methodology. This final chapter brings the thesis to a close while paving the way for further advancements in the field of collision avoidance under non-parametric uncertainty.

Chapter 2

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

2.1 Problem Formulation

We present a formulation for one-step reactive navigation in uncertain environments as a chance-constrained optimization problem (CCO):

$$\min_{\mathbf{u}_t} w_1 \|\overline{\mathbf{v}}_t - \mathbf{v}_d\|_2^2 + w_2 \overline{\mathbf{u}}_t{}^2 \tag{2.1a}$$

$$\Pr(f(\mathbf{x}_t, \theta_t, \mathbf{u}_t, \mathbf{x}_{o,t}, \mathbf{v}_{o,t}) \leq 0) \geq \eta,\ \forall j,\ \mathbf{u}_t \in \mathcal{C} \tag{2.1b}$$

$$\overline{\mathbf{v}}_t = \begin{bmatrix} \overline{v}_t \cos(\overline{\theta}_t + \overline{\omega}_t \Delta t) \\ \overline{v}_t \sin(\overline{\theta}_t + \overline{\omega}_t \Delta t) \end{bmatrix}, \tag{2.2}$$

The cost function (2.1a) involves two key terms. The first term ensures alignment of the nominal velocity with a desired velocity vector $\mathbf{v}_d$, which is typically designed to induce movement towards the goal [17]. Additionally, a control input regularizer is included in the cost function (2.1a). The user-defined weights w_1 and w_2 are used to strike a balance between each cost term. The set $\mathcal{C}$ represents the collection of feasible control inputs, assumed to be convex and formed by affine constraints on $\overline{v}_t, \overline{\omega}_t$. Meanwhile, the inequalities (2.1b) correspond to the so-called chance constraints [8], guaranteeing that the probability of satisfying the velocity obstacle (VO) constraints is greater than or equal to a specified threshold η. The standard form of chance constraints can be seen in the FIg 2.1

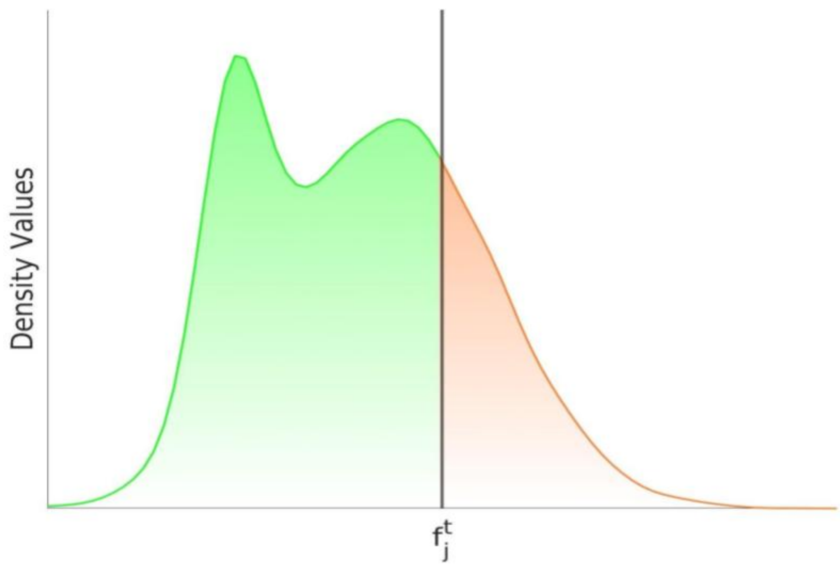

Figure 2.1: Plot depicting the chance constraints. The green part shows the samples which satisfy the constraints, whereas the orange part shows the samples with constraint violation

Solving (2.1a)-(2.1b) presents the main computational challenge, primarily due to the presence of chance constraints. Existing research has largely focused on replacing (2.1b) with alternative, computationally tractable options. A particular approach is discussed in Section V. Notably, many of the existing reformulations assume that the underlying uncertainty follows a Gaussian distribution [4]. In contrast, our objective in this paper is to analyze the impact of Gaussian approximation. Consequently, we introduce our reactive planner, capable of operating with arbitrary uncertainty distributions, ensuring a comprehensive investigation of the effects of different uncertainty models.

2.2 Reformulation as a Distribution Matching Problem

At an intuitive level CCO (2.1a)-(2.1b) has the following interpretation [8]. 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.

$$h(\mathbf{x}_t, \theta_t, \mathbf{u}_t, \mathbf{x}_{o,t}, \mathbf{v}_{o,t}) = \max(0, f(\mathbf{x}_t, \theta_t, \mathbf{u}_t, \mathbf{x}_{o,t}, \mathbf{v}_{o,t})) \tag{2.3}$$

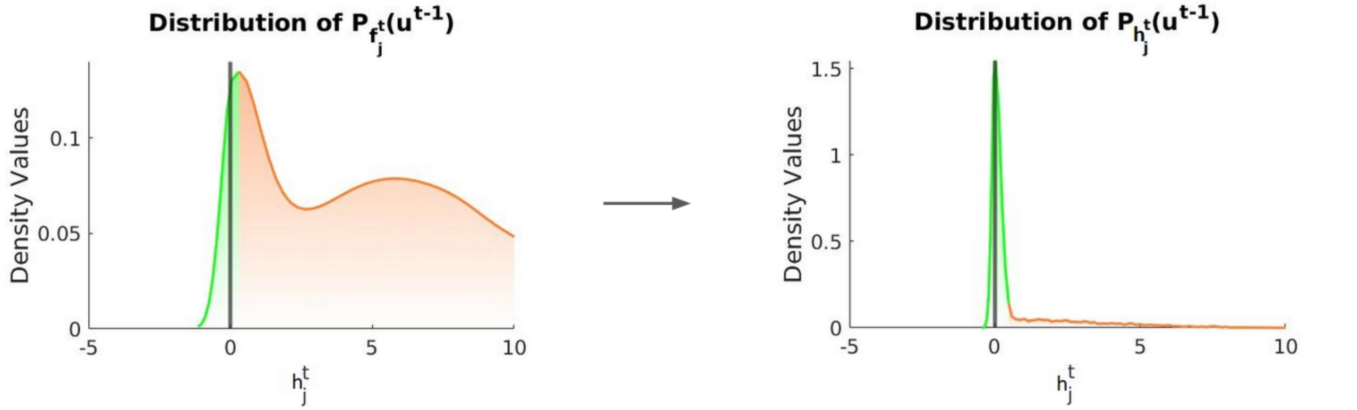

Figure 2.2: These plots depict distribution matching with the Dirac-Delta distribution. The green area represents the samples which are coinciding with the Dirac-Delta distribution. The orange part represent the remaining samples

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 (2.3), 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 (2.1a)-(2.1b) in the following manner:

$$\min_{\overline{\mathbf{u}}_t} l_{dist}(p_h, p_\delta) + w_1\|\overline{\mathbf{v}}_t - \mathbf{v}_d\|_2^2 + w_2{\overline{\mathbf{u}}_t}^2 \tag{2.4}$$

$$\mathbf{u}_t \in \mathcal{C}, \tag{2.5}$$

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.

$$l_{dist}(\mathbf{u}_t) = \overbrace{\|\mu_{p_h}(\mathbf{u}_t) - \mu_{p_\delta}\|_2^2}^{MMD}, \tag{2.6}$$

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

We solve (2.4)-2.5 through a simple control sampling approach. We draw several samples of $\mathbf{u}_t$ from a uniform distribution and then evaluate the cost (2.4) on them. Subsequently, we choose the sample corresponding to the lowest cost. This control input tries to match the constraint violation distribution with the Dirac-Delta distribution as seen in Fig 2.2. Our control sampling relies on efficient evaluation of MMD term to retain online performance. Thus, in the next section, we show how MMD evaluation for a given $\mathbf{u}_t$ can be reduced to computing matrix-matrix products.

2.2.1 Matrix Representation for MMD

The algebraic expression for μ_{p_h} can be derived in the following manner

$$\mu_{p_h} = \sum_{i=0}^{i=N}\sum_{j=0}^{j=N} \alpha_i \beta_j k(\mathbf{h}_{ij}, .) \tag{2.7}$$

where

$$h_{ij} = h(\mathbf{x}_t^i, \theta_t^i, \mathbf{u}_t, \mathbf{x}_{o,t}^j, \mathbf{v}_{o,t}^j) \tag{2.8}$$

and α_i, β_j are constants. Typically, if we draw I.I.D samples, then we have $\alpha_i = \beta_j = \frac{1}{n}$. However, as shown in [8], [7], these constants can be chosen in a clever way to re-weight the importance of each samples leading to sample efficiency. The function $k(.,.)$ is the so-called kernel operator, which in our implementation as Radial Basis Function. That is, $k(\mathbf{c}_1, \mathbf{c}_2) = -\gamma\|\mathbf{c}_1 - \mathbf{c}_2\|_2^2$ for some arbitrary vectors $\mathbf{c}_1, \mathbf{c}_2$.

As clear, μ_{p_h} is formed by first drawing n samples each of robot position/heading $(\mathbf{x}_t^i, \theta_t^i)$ and obstacle position/velocity $((\mathbf{x}_{o,t}^j, \mathbf{v}_{o,t}^j))$ distribution and then evaluating $h(.)$ over all the possible sample pairs. The function $k(.)$ represents the feature map associated with the RBF kernel. We can represent (2.7) in the following more compact form, wherein a_p denotes the p^{th} element of the vector **a**

$$\mu_{p_h} = \sum_{p=1}^{p=N^2} a_p k(\mathbf{h}_p, .), \mathbf{h}_p = \begin{bmatrix} h_{11} \\ h_{12} \\ \vdots \\ h_{nn} \end{bmatrix}, \mathbf{a} = \begin{bmatrix} \alpha_1\beta_1 \\ \alpha_1\beta_2 \\ \vdots \\ \alpha_n\beta_n \end{bmatrix} \tag{2.9}$$

Following a similar approach, we can define μ_{p_δ} as

$$\mu_{p_\delta} = \sum_{q=1}^{q=N^2} b_q k(0, .) \tag{2.10}$$

for some constant vector **b**. Note that (2.10) exploits the fact that the samples from a Dirac-Delta distribution are all zeros.

With respect to the above definition, we can expand (2.6) as

$$\mu_{p_h} - \mu_{P_\delta}{}_2^2 = \mathbf{M}_{cc} - 2\mathbf{M}_{c0} + \mathbf{M}_{00} \tag{2.11a}$$

$$\text{where,} \quad \mathbf{M}_{cc} = \langle \mu_{p_h}, \mu_{p_h} \rangle \tag{2.11b}$$

$$\mathbf{M}_{c0} = \langle \mu_{p_h}, \mu_{P_\delta} \rangle \tag{2.11c}$$

$$\mathbf{M}_{00} = \langle \mu_{P_\delta}, \mu_{P_\delta} \rangle \tag{2.11d}$$

From equation 2.9 and 2.10, we get the following

$$\mathbf{M}_{cc} = \langle \sum_{p=1}^{p=N^2} \mathbf{a}_p k(\mathbf{h}_p, .), \sum_{p=1}^{p=N^2} \mathbf{a}_p k(\mathbf{h}_p, .) \rangle \tag{2.12a}$$

$$\mathbf{M}_{c0} = \langle \sum_{p=1}^{p=N^2} \mathbf{a}_p k(\mathbf{h}_p, .) \sum_{q=1}^{q=N^2} \mathbf{b}_q k(0, .) \rangle \tag{2.12b}$$

$$\mathbf{M}_{00} = \langle \sum_{q=1}^{q=N^2} \mathbf{b}_q k(0, .), \sum_{q=1}^{q=N^2} \mathbf{b}_q k(0, .) \rangle \tag{2.12c}$$

Applying kernel trick on the above equations, we get

$$\mathbf{M}_{cc} = \mathbf{a}_k^T \mathbf{K}_{cc} \mathbf{a}_k \tag{2.13a}$$

$$\mathbf{M}_{c0} = \mathbf{a}_k^T \mathbf{K}_{c0} \mathbf{b}_k \tag{2.13b}$$

$$\mathbf{M}_{00} = \mathbf{b}_k^T \mathbf{K}_{00} \mathbf{b}_k \tag{2.13c}$$

$$l_{dist}(p_h, p_\delta) = \mathbf{a}_k^T \mathbf{K}_{cc} \mathbf{a}_k + \mathbf{a}_k^T \mathbf{K}_{c0} \mathbf{b}_k + \mathbf{b}_k^T \mathbf{K}_{00} \mathbf{b}_k^T \tag{2.13d}$$

$\mathbf{K}_{cc}$, $\mathbf{K}_{c0}$ and $\mathbf{K}_{00}$ are the kernel matrices and are defined as:

$$\mathbf{K}_{cc} = \begin{bmatrix} k(h_{11}, h_{11}) & k(h_{11}, h_{12}) & \dots & k(h_{11}, h_{nn}) \\ k(h_{12}, h_{11}) & k(h_{12}, h_{12}) & \dots & k(h_{12}, h_{nn}) \\ \vdots & \vdots & \vdots & \vdots \\ k(h_{nn}, h_{11}) & k(h_{nn}, h_{12}) & \dots & k(h_{nn}, h_{nn}) \end{bmatrix} \tag{2.14a}$$

$$\mathbf{K}_{c0} = \begin{bmatrix} k(h_{11}, 0) & k(h_{11}, 0) & \dots & k(h_{11}, 0) \\ k(h_{12}, 0) & k(h_{12}, 0) & \dots & k(h_{12}, 0) \\ \vdots & \vdots & \vdots & \vdots \\ k(h_{nn}, 0) & k(h_{nn}, 0) & \dots & k(h_{nn}, 0) \end{bmatrix} \tag{2.14b}$$

$$\mathbf{K}_{00} = 1_{N^2 x N^2} \tag{2.14c}$$

The computation time of evaluating MMD or l_{dist} depends mainly on the computation time of the upper triangle of the symmetric matrix $\mathbf{K}_{cc}$ as the matrix $\mathbf{K}_{00}$ is a set of ones, and $\mathbf{K}_{c0}$ is a column matrix.

Chapter 3

Simulation Results

Implementation Details: All the simulations were carried out on a desktop in Python. The CPU and GPU used were AMD Ryzen 5 3500 and NVIDIA 1660 Super respectively. We queried 100 samples each of robot and obstacle's position and velocity from their distribution to construct the MMD l_{dist} term in optimization (2.4)-(2.5). We reiterate that we don't assume any knowledge on the parametric form for the underlying distribution. We used a fixed set of 625 discrete control inputs to compute the one that led to the lowest value for the cost (2.4). We used $\gamma = 0.1$ in RBF kernel definition. In all the plots demonstrating qualitative results in the form of robot and obstacle trajectories, the blue circle represents the robot's actual position, the yellow circle represents the obstacles' position, and the lighter shade circles surrounding both of them represent the underlying uncertainty in position. Extensive qualitative results and the code can be found at https://github.com/anishgupta31296/MMD-with-Dirac-Delta-Distribution.

Baselines: We call our approach MMD Non-Gaussian when comparing against the following baselines:

- **MMD-Gaussian**: This baseline follows the same approach of distribution matching in RKHS through MMD. The only difference with our approach is that it computes a Gaussian approximation of the motion and perception noise.

- **KLD**: This baseline from [7] also follows the interpretation of CCO as a distribution matching problem. But it differs from our approach in the following respects. First, it fits a Gaussian Mixture Model to the noise distribution. Second, it works with the distribution of VO constraints, while our approach uses the distribution of violations.

- **PVO** : This baseline from [18] proposed a deterministic reformulation of the chance constraints over VO presented in (2.1b). However, it requires computing the Gaussian approximation of motion and perception noise.

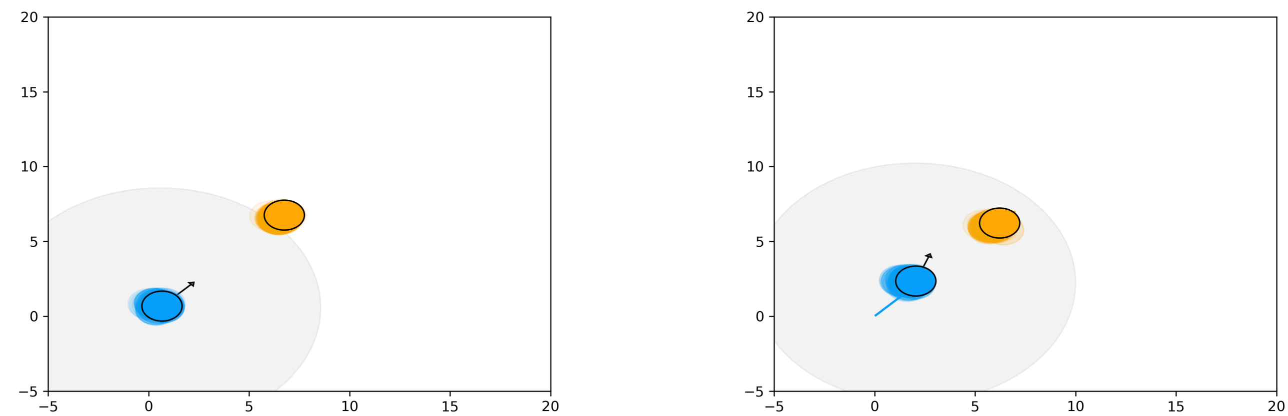

(a) Beginning of collision avoidance. The VO constraint violation distribution is very far from Dira-Delta.

(b) Some part of the VO constraint violation distribution coincides with the Dirac Delta Distribution

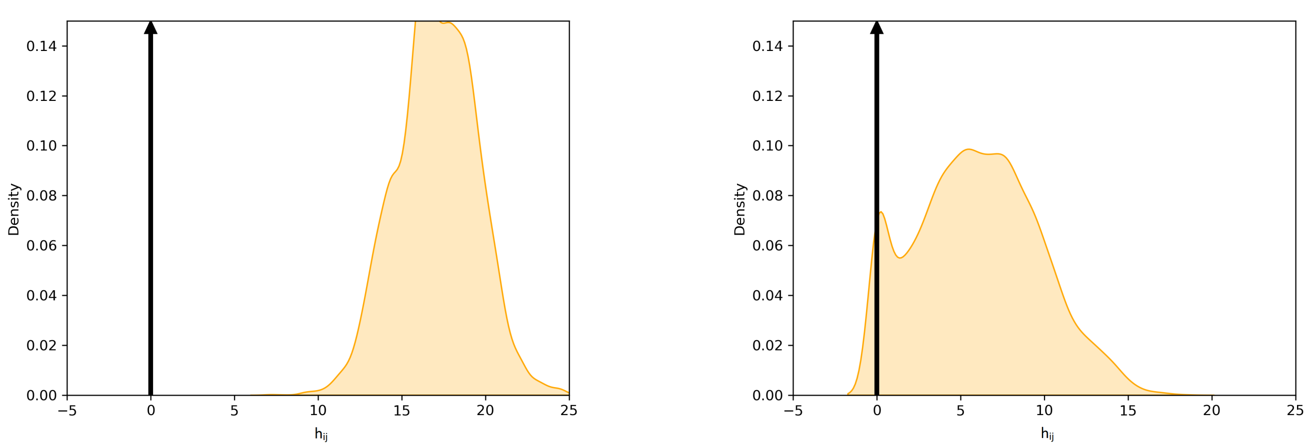

Figure 3.1: Validation of distribution matching interpretation of CCO.

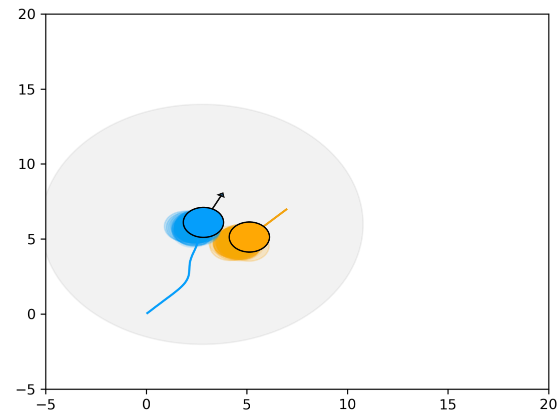

(c) Towards the end of the collision avoidance maneuver, the distribution of the VO constraint violations becomes close to that of the Dirac-Delta

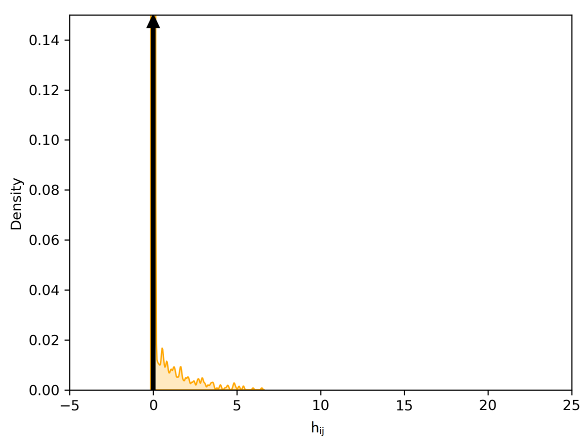

Figure 3.1: Validation of distribution matching interpretation of CCO.

3.0.1 Validating Distribution Matching Interpretation

Fig.3.1 shows a simple scenario where a robot has an imminent head-on collision with an obstacle. Fig.3.1a shows at the start of the collision avoidance maneuver, the distribution of VO constraint violation is entirely on the right of zero. As the robot computes collision avoidance maneuver by solving (2.4)-2.5, the constraint violations (almost) converge to the Dirac-Delta distribution.

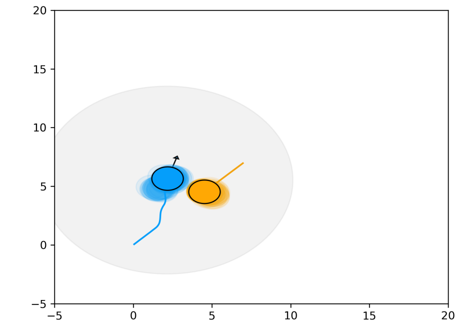

(a) Favorable homotopy chosen by the robot

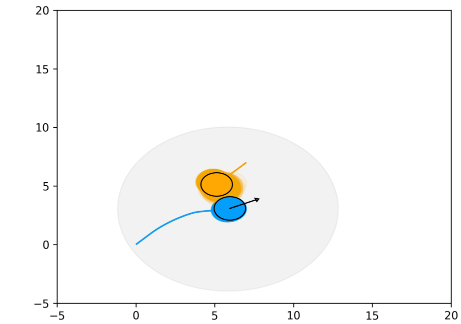

(b) Un-favorable homotopy chosen by the robot

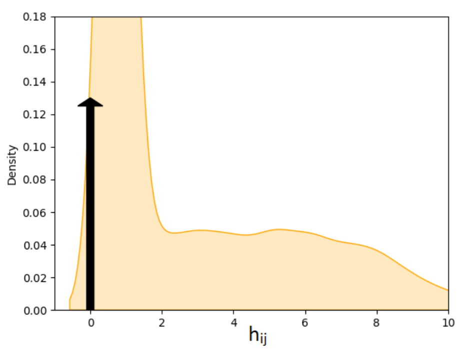

(c) Corresponding VO constraint violation distribution

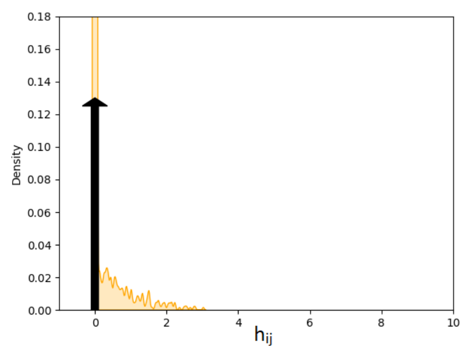

(d) Corresponding VO constraint violation distribution

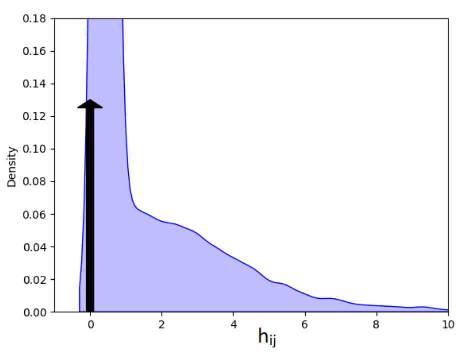

(e) Corresponding VO constraint violation distribution under Gaussian approximation

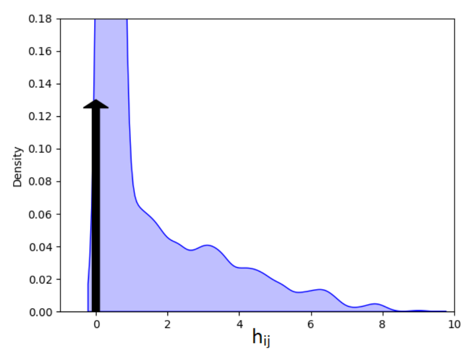

(f) Corresponding VO constraint violation distribution under Gaussian approximation

Figure 3.2: These figures illustrate how better favourable homotopy selection will lead to better distribution matching and hence, larger number of samples will be avoided

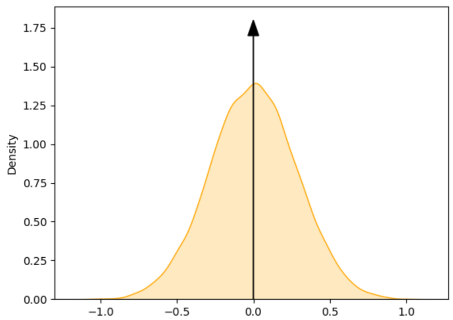

(a) Distribution 1

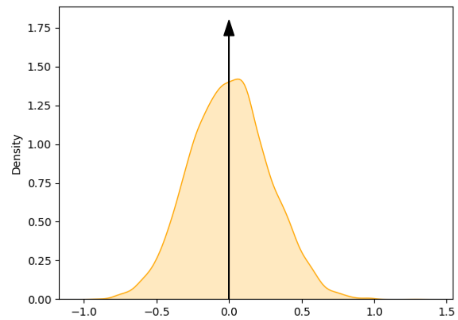

(b) Distribution 2

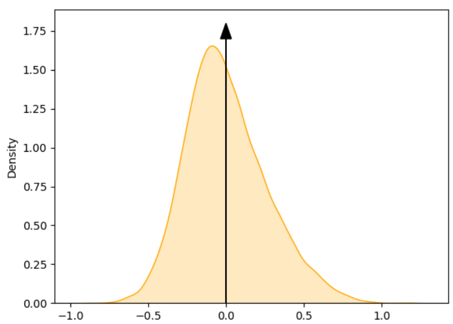

(c) Distribution 3

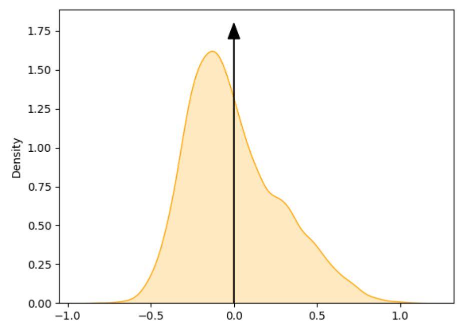

(d) Distribution 4

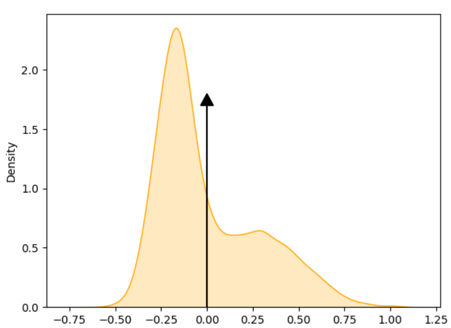

(e) Distribution 5

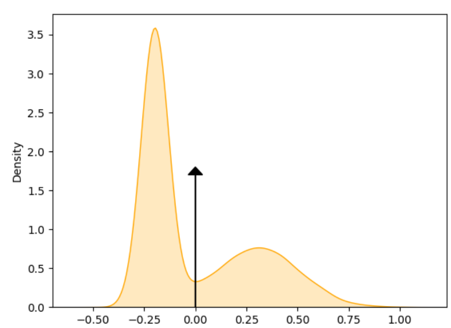

(f) Distribution 6

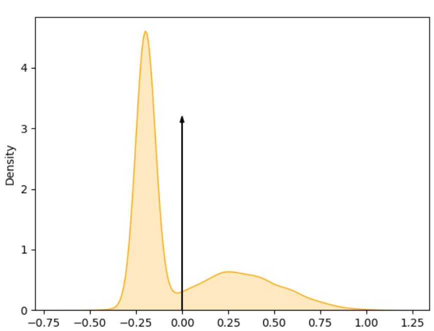

(g) Distribution 7

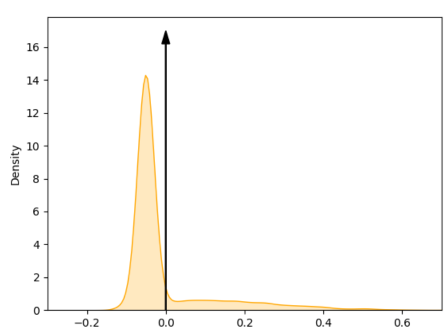

(h) Distribution 8

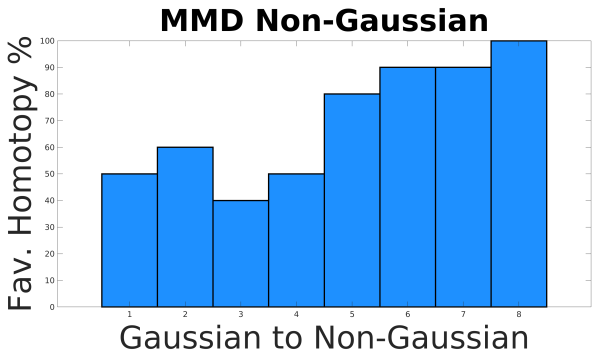

(a) This bar plot depicts the frequency of choosing favorable homotopy with our approach MMD Non-Gaussian in a single obstacle benchmark shown in Fig.3.2c,3.2e under the noise distributions from Figure 3.3. We can observe the increasing likelihood of choosing the favorable homotopy as we move towards non-Gaussian noise distributions

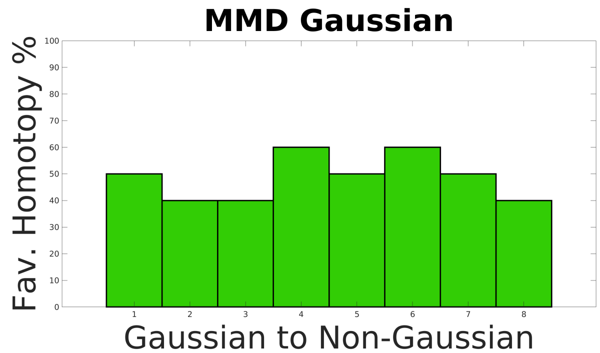

(b) This bar plot depicts the frequency of choosing favourable homotopy for MMD Gaussian in case of single obstacle benchmark shown in Fig.3.2c,3.2e using the noise distributions in Figure 3.3. We can see the probability of choosing the correct side remains similar even when the noise distribution becomes increasingly non-Gaussian.

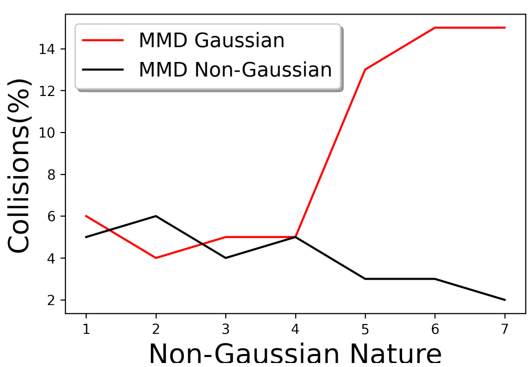

(c) Effect of non-gaussian nature on number of samples colliding

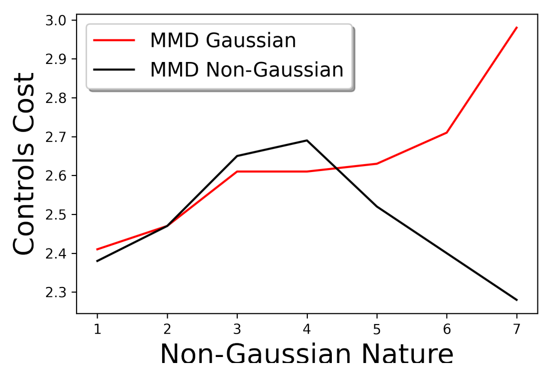

(d) Effect of non-gaussian nature on control costs. The x-axis shows the distribution number from Fig.3.3

Figure 3.4: Quantitative Analysis on non-gaussian nature of distribution. The x-axis shows the distribution number from Fig.3.3

3.0.2 Analyzing the Choice of Homotopies

This section presents the most important empirical result of our paper. We consider a benchmark with a single obstacle as shown in Fig.3.2 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.3.2 which results in the robot passing the obstacle from different sides. Clearly, Fig.3.2a is the favorable homotopy in this scenario while that shown in Fig.3.2b leads to a large overlap between the robot and obstacle position uncertainty. Fig.3.2c 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.3.2d that recreates the constraint violation distribution for the control input leading to unfavorable homotopy. We can clearly see a stark difference between Fig.3.2c and 3.2d. Now, we hypothesize that any planner that can capture the true distribution of constrint violation for a given control input can easily distinguish between a favorable and unfavorable homotopy. We will soon discuss how our MMD Non-Gaussian planner in fact fits this description. But before that, we turn our attention to Fig3.2e and 3.2f that 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. In other words, the Gaussian approximation erroneously has made both homotopies equally bad/good. As a result, it is not possible to reliably distinguish between favorable and unfavorable homotopies.

To further strengthen our claims, we design one more experiment. In Fig.3.3, 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.3.2 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.3.4a shows the percentage of times the

Table 3.1

Method	Computation Time(s)	Success-Rate(%)
MMD Non-Gaussian	0.06	95.5
MMD Gaussian	0.07	72
PVO	0.03	89
KLD(GMM-fit)	0.06(1.44)	82

robot chooses homotopy of Fig.3.2a over that of Fig.3.2b. When the actual noise is Gaussian, the robot randomly chooses either homotopy. In fact for Gaussian noise, there is no real benefit provided by one homotopy over another. 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.3.4b. Under Gaussian approximation, the robot always chooses the homotopies randomly.

Fig.3.4c and 3.4d 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.

3.0.3 Quantitative Comparisons

In this section, we compare our MMD Non-Gaussian formulation with MMD Gaussian, KLD and PVO baselines defined in the beginning of section 3. The comparisons are shown in the bar plots of Figure 3.6. Fig.3.5 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 can fits a complicated a GMM to the motion and perception noise.

Figure 3.6a 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.3.6b 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 how many of the drawn position samples from the robot uncertainty collide with that of the obstacles for all the baselines. This metric serves as a proxy of collision probability. Our approach consistently maintains the percentage value at 5 or less. All other baselines performance varies over the benchmarks and lies between $11 - 28\%$. This is further reiterated in Table 3.1

Table 3.1 compares the computation time for our approach and all the baselines. The PVO approach is the fastest while the rest of the approaches have comparable run-times.

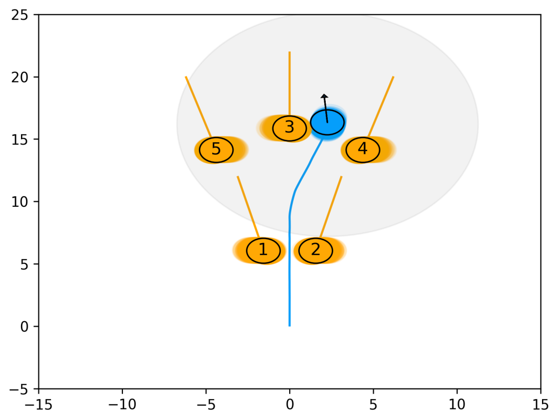

(a) MMD Non-Gaussian

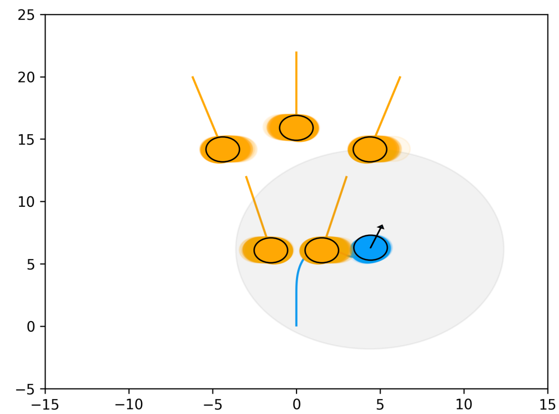

(b) MMD Gaussian

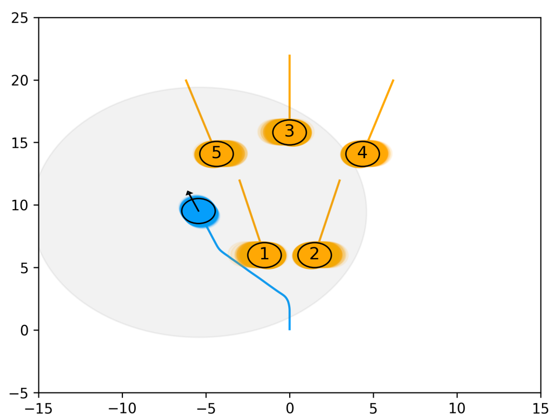

(c) PVO

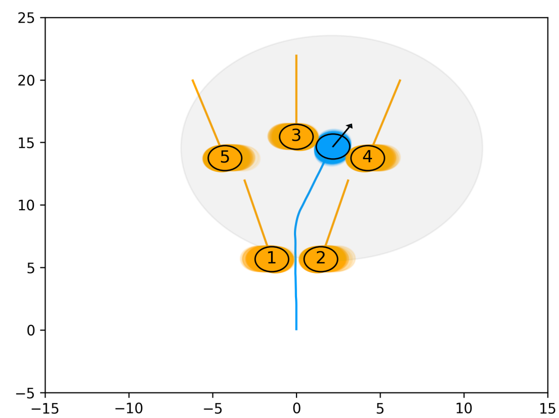

(d) KLD

Figure 3.5: Collision avoidance using MMD Non-Gaussian and various baselines for 5 obstacle case

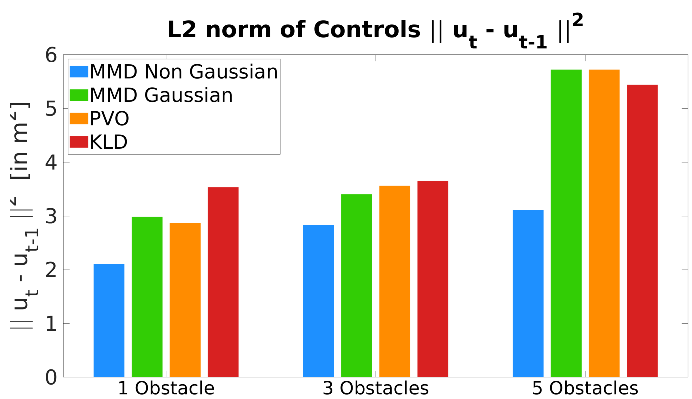

(a) Control Costs comparison

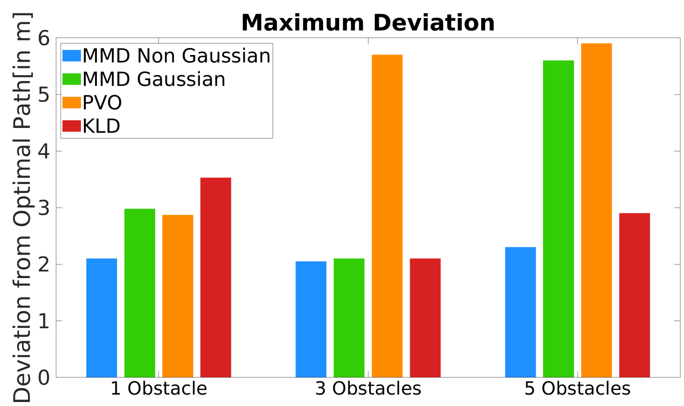

(b) Deviation from optimal path comparison

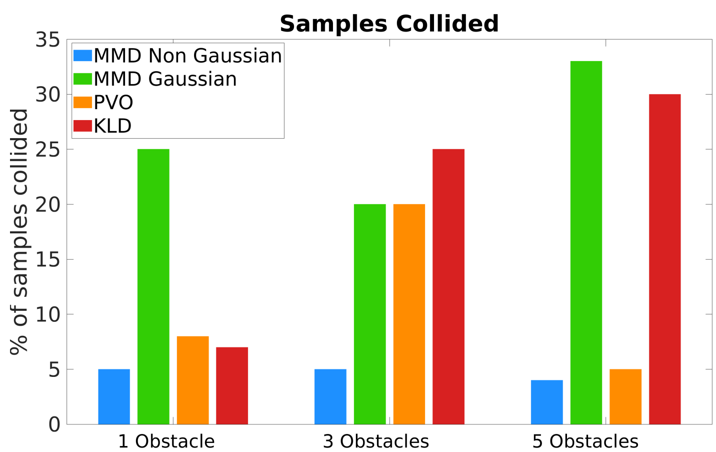

(c) Number of colliding samples

Figure 3.6: 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.

Chapter 4

Conclusion and Future Work

This thesis extensively utilizes the distribution matching interpretation of chance-constrained optimization (CCO) to propose a novel approach for reactive dynamic obstacle avoidance. The central focus is on minimizing the deviation of the distribution of constraint violations from Dirac-Delta, enabling a detailed analysis of how bias in non-Gaussian motion and perception noise can be leveraged to strategically choose favorable homotopies for efficient collision avoidance. Moreover, a crucial aspect of this research is the investigation of the Gaussian approximation of uncertainty, which tends to treat all homotopies as equally good or bad, ultimately compelling the planner to select sub-optimal motions. By revealing this limitation, our study emphasizes the importance of adopting non-parametric models to achieve more accurate and reliable collision avoidance strategies.

This thesis marks the first presentation of a comprehensive analysis of non-Gaussian motion and perception noise and its influence on homotopy selection for collision avoidance. The insights gained from this research hold significant implications for advancing autonomous robotic systems in uncertain environments. Furthermore, our work not only uncovers the potential of bias in non-parametric distributions but also identifies the drawbacks of Gaussian approximations, shedding light on crucial considerations for future robotic motion planning.

Building on the foundation laid by this thesis, our future research endeavors will extend the proposed reactive approach into a full-fledged multi-step Model Predictive Control (MPC) setting. This strategic expansion will enable us to address complex and multi-dimensional planning scenarios, enhancing the overall performance and adaptability of our approach in real-world dynamic environments.

In summary, this thesis presents a groundbreaking exploration into the realm of collision avoidance under non-parametric uncertainty, unveiling the potential of bias in non-Gaussian distributions for more effective homotopy selection. The empirical findings, computational analyses, and theoretical contributions presented herein contribute significantly to the advancement of autonomous robotics and pave the way for further developments in motion planning in challenging and uncertain conditions.

Related Publications

1. **Leveraging Distributional Bias For Reactive Collision Avoidance under Uncertainty: A Kernel Embedding Approach**
   Anish Gupta*, Arun Kumar Singh2, and K. Madhava Krishna
   2022 IEEE 18th International Conference on Automation Science and Engineering (CASE)

Other Publications

1. **Multi-modal model predictive control through batch non-holonomic trajectory optimization: Application to highway driving**
   Vivek K Adajania, Aditya Sharma, **Anish Gupta**, Houman Masnavi, K Madhava Krishna, Arun K Singh
   IEEE Robotics and Automation Letters (Volume: 7, Issue: 2, April 2022) Page(s): 4220 - 4227

2. **Non Holonomic Collision Avoidance under Non-Parametric Uncertainty: A Hilbert Space Approach**
   Unni Krishnan R Nair*, **Anish Gupta***, D. A. Sasi Kiran, Ajay Shrihari, Vanshil Shah, Arun Kumar Singh, K. Madhava Krishna
   2021 European Control Conference (ECC)

Bibliography

[1] Hai Zhu and Javier Alonso-Mora. Chance-constrained collision avoidance for mavs in dynamic environments. *IEEE Robotics and Automation Letters*, 4(2):776–783, 2019.

[2] Jae Sung Park and Dinesh Manocha. Efficient probabilistic collision detection for non-gaussian noise distributions. *IEEE Robotics and Automation Letters*, 5(2):1024–1031, 2020.

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

[4] Hai Zhu and Javier Alonso-Mora. Chance-constrained collision avoidance for mavs in dynamic environments. *IEEE Robotics and Automation Letters*, 4(2):776–783, 2019.

[5] David M Rosen and John 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.

[6] Anirudha Majumdar and Vincent Pacelli. Fundamental performance limits for sensor-based robot control and policy learning. *arXiv preprint arXiv:2202.00129*, 2022.

[7] SriSai Naga Jyotish Poonganam, Bharath Gopalakrishnan, Venkata Seetharama Sai Bhargav Kumar Avula, Arun Kumar Singh, K. Madhava Krishna, and Dinesh Manocha. Reactive navigation under non-parametric uncertainty through hilbert space embedding of probabilistic velocity obstacles. *IEEE Robotics and Automation Letters*, 5(2):2690–2697, 2020.

[8] Bharath Gopalakrishnan, Arun Kumar Singh, K Madhava Krishna, and Dinesh Manocha. Solving chance-constrained optimization under nonparametric uncertainty through hilbert space embedding. *IEEE Transactions on Control Systems Technology*, 2021.

[9] Paolo Fiorini and Zvi Shiller. Motion planning in dynamic environments using velocity obstacles. *The International Journal of Robotics Research*, 17(7):760–772, 1998.

[10] Bharath Gopalakrishnan, Arun Kumar Singh, and K.Madhava 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)*, pages 4961–4968, 2015.

[11] Stephen Boyd and Almir Mutapcic. Stochastic subgradient methods. *Lecture Notes for EE364b, Stanford University*, 2008.

[12] Lars Blackmore, Masahiro Ono, Askar Bektassov, and Brian C Williams. A probabilistic particle-control approximation of chance-constrained stochastic predictive control. *IEEE transactions on Robotics*, 26(3):502–517, 2010.

[13] Ali Mesbah, Stefan Streif, Rolf Findeisen, and Richard D Braatz. Stochastic nonlinear model predictive control with probabilistic constraints. In *2014 American control conference*, pages 2413–2419. IEEE, 2014.

[14] Brandon Luders, Mangal Kothari, and Jonathan How. Chance constrained rrt for probabilistic robustness to environmental uncertainty. In *AIAA guidance, navigation, and control conference*, page 8160, 2010.

[15] Bruce G Lindsay and Prasanta Basak. Moments determine the tail of a distribution (but not much else). *The American Statistician*, 54(4):248–251, 2000.

[16] Sudarshan S Harithas, Rishabh Dev Yadav, Deepak Singh, Arun Kumar Singh, and K Madhava Krishna. Cco-voxel: Chance constrained optimization over uncertain voxel-grid representation for safe trajectory planning. *arXiv preprint arXiv:2110.02904*, 2021.

[17] 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*, pages 1928–1935, May 2008.

[18] Bharath Gopalakrishnan, Arun Kumar Singh, and K Madhava Krishna. Closed form characterization of collision free velocities and confidence bounds for non-holonomic robots in uncertain dynamic environments. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pages 4961–4968. IEEE, 2015.