

MMD-OPT: Maximum Mean Discrepancy-Based Sample Efficient Collision Risk Minimization for Autonomous Driving

Basant Sharma¹, Graduate Student Member, IEEE, and Arun Kumar Singh

Abstract—We propose MMD-OPT: a sample-efficient approach for minimizing the risk of collision under arbitrary prediction distribution of the dynamic obstacles. MMD-OPT is based on embedding distribution in Reproducing Kernel Hilbert Space (RKHS) and the associated Maximum Mean Discrepancy (MMD). We show how these two concepts can be used to define a sample efficient surrogate for collision risk estimate. We perform extensive simulations to validate the effectiveness of MMD-OPT on both synthetic and real-world datasets. Importantly, we show that trajectory optimization with our MMD-based collision risk surrogate leads to safer trajectories at low sample regimes than popular alternatives based on Conditional Value at Risk (CVaR).

Note to Practitioners—Autonomous Driving software stacks have dedicated modules for predicting trajectories of the obstacles (neighboring vehicles). Typically, these predictors provide a set of possible future motions for the obstacles, each of which can have different likelihoods of happening. Thus, a key challenge is to reason about collision risk in a given scene based on predicted trajectories, but without being overly conservative. For example, treating each predicted trajectory as a separate obstacle, without any attention to their likelihood, may not allow any feasible motion to the ego-vehicle. Our work addresses this challenge by proposing a probabilistic approach for modeling and minimizing collision risk. Our core impact lies in improving sample efficiency: that is assessing and minimizing collision risk based on just a handful of predicted trajectories for a given obstacle. Our approach can be easily integrated with any deep neural network based trajectory predictors which have become the de facto standard in autonomous driving industry. Our formulation easily extends to applications like indoor navigation with mobile robots, since human trajectory predictors have structural similarity with those deployed in autonomous driving. A practical limitation of our approach is that it requires additional computing power (in the form of GPU accelerators) to achieve real-time performance.

Index Terms—Autonomous driving, planning under uncertainty, collision risk, obstacle avoidance, motion planning.

I. INTRODUCTION

ENSURING collision avoidance between the ego vehicle and the dynamic (neighbouring vehicles) obstacles is a

Received 2 November 2024; revised 18 April 2025; accepted 11 July 2025. Date of publication 18 July 2025; date of current version 1 August 2025. This article was recommended for publication by Associate Editor B. Lacevic and Editor J. Yi upon evaluation of the reviewers' comments. This work was supported in part by the Estonian Research Council under Grant PSG753, in part by the Bolt Technologies under Project LLTAT21278, and in part by European Union and Estonian Research Council under Project TEM-TA101. (Corresponding author: Basant Sharma.)

The authors are with the Institute of Technology, University of Tartu, 50090 Tartu, Estonia (e-mail: basant.sharma@ut.ee; arun.singh@ut.ee).

Data is available on-line at <https://github.com/Basant1861/MMD-OPT>. This article has supplementary downloadable material available at <https://doi.org/10.1109/TASE.2025.3589941>, provided by the authors.

Digital Object Identifier 10.1109/TASE.2025.3589941

fundamental requirement in autonomous driving. This in turn, requires predicting the trajectories of the latter over a time horizon. It is often difficult to predict just one deterministic trajectory for the dynamic obstacles because of the uncertainty associated with its intentions and the deployed sensors. Thus, existing works [1], [2], [3] are focused on predicting a distribution of trajectories that can capture their various likely motions.

The objective of this paper is to leverage the predictions from off-the-shelf trajectory predictors in the best possible manner for safe motion planning. Specifically, the core emphasis is on developing a risk cost that captures the probability of collision given trajectory predictions of the dynamic obstacles. This collision risk cost can then be minimized with any off-the-shelf optimizer along with other costs associated with control, smoothness and lane adherence.

A. Existing Gaps

There are two core challenges with regards to developing a collision risk cost. First, the analytical form of prediction distribution may be intractable or unknown. For example, the state-of-the-art trajectory predictors like [1], [4], [5], [6] characterize the predicted trajectory distribution in the form of deep generative models that can be arbitrary complex with multiple modes. Thus, existing collision risk derived from chance-constraints under Gaussian uncertainty [7], [8], [9], [10] become unsuitable.

Second, due to the complex nature of the predicted distribution, we often just have access to samples drawn from it. In such a case, the collision risk cost should be amenable to sample efficient estimation. That is, it should capture the true probability of collision with just a few samples of predicted obstacle trajectories. The sample efficiency is also directly related to computational run-time. The use of higher number of predicted trajectory samples translates to an increase in the number of collision checks between the ego and the obstacles.

B. Contributions

In this paper, we propose MMD-OPT that has two novel components: (i) a sample efficient surrogate for collision risk cost that can work with arbitrary prediction distribution of obstacle motions, (ii) a trajectory optimizer for minimizing the developed surrogate cost.




Fig. 1. The figure shows a scenario where an obstacle has multiple intents (lane-change vs lane-following), each associated with a trajectory distribution. However, both intents have wildly different probabilities. In this particular example, the probability of lane-change is higher. For safe navigation, the ego-vehicle needs to consider this multi-modal nature of obstacle trajectories while planning its own motions. Our proposed approach estimates the more likely samples (the reduced-set) from a set of obstacle trajectories sampled from an unknown distribution. This allows us to plan probabilistically safe motions while appropriately discriminating the low and high-probability obstacle manoeuvres.

The foundation of our work is built around estimating how different two given distributions are. For example, in the case of Gaussian, this can be calculated using the Kullback-Liebler (KL) divergence. However, computing the dissimilarity between arbitrary distributions is often intractable. One possible workaround is to compute the difference between two distributions in the Reproducing Kernel Hilbert Space (RKHS) [11], [12]. Specifically, given only drawn samples, we can represent the underlying distribution as a point in RKHS. Moreover, the difference between the two embedding can be easily captured by the so-called Maximum Mean Discrepancy (MMD) measure.

1) *Algorithmic Contribution*: MMD-OPT uses the concept of RKHS embedding and MMD in two novel ways which also forms our core algorithmic contribution. First, given a set of predicted obstacle trajectories, we embed the underlying distribution of collision constraint residuals in RKHS. Subsequently we define a collision risk surrogate as the MMD between the RKHS embedding of collision constraint residuals and Dirac-Delta distribution. The advantage of our MMD-based surrogate is that we can systematically improve its sample efficiency by leveraging the underlying properties of RKHS embedding and MMD. To this end, we propose a bi-level optimization problem that tells us which of the obstacle trajectory samples can be disregarded without compromising the ability of our MMD-based surrogate to capture the true probability of collision (see Fig. 1). This in turn, reduces the number of collision checks required to reliably produce safe trajectories and adds to the overall sample and computational efficiency of MMD-OPT. Finally, we propose a custom sampling-based optimizer to minimize MMD-based collision risk surrogate along with driving discomfort in the form of sharp accelerations and lane boundary violations.

2) *State-of-the-Art Performance*: We empirically compare planning with our MMD-based collision risk surrogate with




Fig. 2. The image shows that nonlinear reference paths (such as curvy roads) are represented by straight lines on the s -axis in Frenet coordinates. However, motions that do not exactly follow the reference path result in nonlinear motions in Frenet coordinates. Instead, such motions result in an offset from the reference path and therefore the s -axis, which is described with the d coordinate.

that achieved using popular risk-cost alternatives derived from Sample Average Approximation (SAA) [13] and Conditional Value at Risk (CVaR) [14]. We show that our approach leads to safer trajectories for a given number of samples of obstacle trajectories. The difference is more prominent when the underlying prediction distribution departs significantly from the Gaussian form. We also demonstrate that our approach is computationally fast enough for real-time applications.

C. Organization of the Paper

The paper is structured as follows: Section II defines the motion model, formulates risk-aware trajectory optimization, and introduces exact and sample-based collision risk approximations. Section III provides an overview of RKHS embeddings, kernel functions, and the Maximum Mean Discrepancy (MMD) metric. Section IV presents MMD-OPT, including an MMD-based risk surrogate, reduced-set selection, bi-level optimization, and a sampling-based trajectory optimizer. Section V discusses connections to chance-constrained optimization, risk-aware planning, and RKHS applications. Section VI details implementation, evaluates performance in static and dynamic settings, compares MMD-OPT with CVaR and SAA, and analyzes scalability. Section VII concludes with key findings and future research directions.

II. PROBLEM FORMULATION AND PRELIMINARIES

Symbols and Notations Scalars will be represented by normal-font small-case letters, while bold-faced variants will represent vectors. We will use upper-case bold fonts to represent matrices. Symbol k will represent the discrete-time step while T will represent the transpose operator.

A. Motion Model in Frenet Frame

We formulate trajectory planning of the ego-vehicle in the road-aligned reference frame known as the Frenet frame (Fig 2). In this setting, the curved roads can be treated as ones with straight-line geometry. This is achieved by aligning the longitudinal and lateral motions of the ego-vehicle with the axes of the Frenet frame. Specifically, in Frenet coordinates,

we use two variables, namely, s representing longitudinal displacement or the distance along the reference path and d representing the lateral displacement or perpendicular offset from the reference path. The s coordinate starts with $s = 0$ at the beginning of the reference path. Lateral coordinates on the reference path are represented with $d = 0$. Depending on the convention used for the Frenet frame, d is positive to the left or right of the reference path.

We adopt the following bicycle model in the Frenet frame for the ego-vehicle [15], where in the subscript k represents the time-step.

$$s_{k+1} = s_k + \dot{s}_k \Delta t, d_{k+1} = d_k + \dot{d}_k \Delta t \quad (1a)$$

$$\psi_{k+1} = \psi_k + \dot{\psi}_k \Delta t, v_{k+1} = v_k + a_k \Delta t \quad (1b)$$

$$\dot{s}_{k+1} = \frac{v_k \cos \psi_k}{1 - d_k \kappa(s_k)}, \dot{d}_{k+1} = v_k \sin \psi_k \quad (1c)$$

$$\dot{\psi}_{k+1} = \frac{v_k \tan(\theta_k)}{B} - \kappa(s_k) \frac{v_k \cos \psi_k}{1 - d_k \kappa(s_k)} \quad (1d)$$

where ψ_k is the heading of the ego-vehicle in the Frenet Frame and B is the wheelbase of the car. The ego-vehicle is driven by a combination of longitudinal acceleration a_k and steering input θ_k . The variable $\kappa(s_k)$ is the curvature of the reference path at the longitudinal coordinate s_k .

1) *Differential Flatness*: The motion model (1a)-(1d) is differentially flat. That is, we can represent the control inputs as a function of position and its derivatives. More precisely, we have

$$v_k = \sqrt{(\dot{s}_k(1 - d_k \kappa(s_k)))^2 + (\dot{d}_k)^2}, a_k = \frac{v_{k+1} - v_k}{\Delta t} \quad (2)$$

$$\psi_k = \tan^{-1}(\dot{d}_k, \dot{s}_k), \dot{\psi}_k = \frac{\psi_{k+1} - \psi_k}{\Delta t} \quad (3)$$

$$\theta_k = \tan^{-1} \left(\frac{(\dot{\psi}_k + \kappa(s_k) \dot{s}_k) B}{v_k} \right) \quad (4)$$

Remark 1: Eqns.(2)-(4) is leveraged later in the paper (Eqn. (32)) to express control costs over θ_k and a_k in terms of position-level trajectory (s_k, d_k) and their derivatives.

B. Risk-Aware Trajectory Optimization

Let τ_k represent the obstacle position at time-step k . Let $\tau \in \mathbb{R}^{2H}$ denote the obstacle trajectory formed by stacking time-stamped way-points over a planning horizon H . In the stochastic case, τ is a random variable with distribution p_τ . With these notations in place, we can leverage differential flatness to formulate risk-aware trajectory optimization directly in the trajectory space as shown below:

$$\min_{\mathbf{s}, \mathbf{d}} c(\mathbf{s}^{(q)}, \mathbf{d}^{(q)}) + r(\mathbf{s}, \mathbf{d}, \tau) \quad (5)$$

$$\text{s.t. } \mathbf{h}(\mathbf{s}^{(q)}, \mathbf{d}^{(q)}) = \mathbf{0} \quad (6)$$

$$\mathbf{g}(\mathbf{s}^{(q)}, \mathbf{d}^{(q)}) \leq \mathbf{0}, \forall k$$

$$\begin{aligned} \mathbf{s}^{(q)} &= (s_1^{(q)}, s_2^{(q)}, \dots, s_k^{(q)}, \dots, s_H^{(q)}) \\ \mathbf{d}^{(q)} &= (d_1^{(q)}, d_2^{(q)}, \dots, d_k^{(q)}, \dots, d_H^{(q)}), \quad q \in \{0, 1, 2\} \end{aligned} \quad (7)$$

The first term in the cost (5) minimizes some function $c(\cdot)$ defined over the q^{th} derivative of the position variable. We can use the differential flatness property to roll any state and control dependent cost into this term (e.g. see (32)).

The second term, $r(\cdot)$ models the collision risk of the ego-vehicle colliding with the obstacle following the predicted trajectory τ . Since obstacles move unpredictably, τ represents a random variable, which can attain many potential instantiations. The risk cost $r(\cdot)$ evaluates the likelihood of the ego-vehicle's planned motion intersecting with any of the potential instantiations of τ . The equality constraints enforce the initial and final boundary conditions on the planned ego trajectory. The inequality constraints (7) model the different kinematic and geometric constraints (e.g. lane boundary) that need to be satisfied by the optimal trajectory. We provide the exact mathematical form for these constraints in Section VI.

Remark 2: For ease of exposition, the problem formulation (5)-(7) only considers the case of a single neighbouring vehicle that acts as an obstacle to the ego-vehicle. Thus, we only have one random obstacle trajectory τ . The extension to arbitrary traffic density is trivial.

Remark 3: The Optimization (5)-(7) is solved under the assumption that the probability distribution p_τ of τ can be arbitrary. But we can draw samples from p_τ . Moreover, sampling is relatively cheap compared to the computational cost of estimating the risk r (defined below) from the predicted samples of obstacle trajectories.

Remark 4: The Optimization (5)-(7) represents a single instance planning problem. We can use the optimizer in a receding horizon a.k.a the model predictive control setting by solving (5)-(7) from the current state at each control cycle.

C. Exact Collision Risk

Let $f_k(s_k, d_k, \tau_k)$ represent the collision constraint function at time step k , such that $f_k \leq 0, \forall k$ ensures collision avoidance. Then, we can define the worst-case collision-constraint value (f) in the following manner.

$$f(\mathbf{s}, \mathbf{d}, \tau) = \max_k (f_k(s_k, d_k, \tau_k)) \quad (8)$$

where we recall that $\tau_k = (s_{o,k}, d_{o,k})$ is the position of the obstacle at time-step k and is obtained from the trajectory predictor. Stacking τ_k at different time steps gives us τ . There are several ways to define $f_k(s_k, d_k, \tau_k)$. In the most general setting, this can be an unknown function derived from 3D occupancy maps [16]. We can also derive an analytical form for $f_k(s_k, d_k, \tau_k)$ if we approximate the shape of the ego vehicle and the neighboring vehicles in terms of geometric primitives such as ellipses.

In the stochastic setting, when τ is a random variable, f provides a distribution of the worst-case collision-constraint value. Thus, using (8), we can define collision risk through (9), wherein P denotes the probability.

$$r = P(f(\mathbf{s}, \mathbf{d}, \tau) \geq 0). \quad (9)$$

The r.h.s of (9) maps the ego and the obstacle trajectory to the probability that the worst-case collision-constraint value is greater than zero and thus, indicating an imminent collision. Thus, minimizing r within (5)-(7) will produce safer trajectories. However, the main computational challenge stems from the fact that for arbitrary obstacle trajectory distribution, p_τ , the r.h.s of (9) does not even have an analytical




Fig. 3. Figure shows the distribution (p_f) of worst-case collision constraint value (f). The distribution ($p_{\bar{f}}$) shows the distribution of collision-constraint residuals. Note that by definition, the mass of $p_{\bar{f}}$ lies to the right of $\bar{f} = 0$. We want the mass of p_f to lie completely to the left of $f = 0$. Thus, the risk-aware trajectory optimization is essentially finding the ego-vehicle trajectory that reshapes the distribution of p_f in appropriate form. Moreover, as the mass of p_f shifts to the left, the distribution $p_{\bar{f}}$ converges to a Dirac-Delta distribution.

expression. This necessitates the development of approximations/surrogates that can achieve a similar effect as r while being computationally more tractable. This is discussed next.

D. Sample Approximation of Risk Metric

In this sub-section, we will use the samples of obstacle trajectories to approximate r . We begin by defining the collision constraint residual function \bar{f} in the following manner

$$\bar{f} = \max(0, f(\mathbf{s}, \mathbf{d}, \tau)) \quad (10)$$

The \bar{f} as defined in (10) maps the distribution over ego and obstacle trajectories to a distribution over collision constraint residuals. The relationship between the probability distribution of f and \bar{f} is shown in Fig. 3. Now, assume that we have access to N samples of obstacle trajectories ${}^j\tau, \forall j = 1, 2, \dots, N$ drawn from p_τ . The different approximations of r revolve around using the obstacle trajectories to obtain samples of \bar{f} and computing appropriate empirical statistics on them.

1) *Sample Average Approximation (SAA) Risk Estimate:* A simple yet effective approximation of collision risk r (9) can be obtained by simply averaging the collision constraint residuals across different samples of the predicted trajectories [13], [17]. That is we have

$$r_{SAA} = \frac{1}{N} \sum_{j=1}^N \mathbb{I}({}^j\bar{f}(\mathbf{s}, \mathbf{d}, {}^j\tau) > 0) \quad (11)$$

where $\mathbb{I}(\cdot)$ is an indicator function returning 1 when the conditions in the operand are satisfied and 0 otherwise. In essence, (11) minimizes the frequency of constraint violation over the samples of predicted obstacle trajectories. Although conceptually simple, estimates of the form (11) have proved useful in the existing literature [13].

2) *Conditional Value at Risk (CVaR) Based Risk Estimate:* An alternative sample approximation of collision risk can be obtained by computing the empirical/sample estimate of CVaR




Fig. 4. Kernel mean embedding of probability distributions p_δ and $p_{\bar{f}}$ into RKHS \mathcal{H} .

of \bar{f} , using the samples $\{\bar{f}^j\}_{j=1}^{j=N}$ [14], [18], [19]. That is we define the risk metric as

$$r_{CVaR}^{emp} = CVaR_\alpha^{emp}(\bar{f}) \quad (12)$$

where α is some tuneable parameter associated with CVaR. We can follow the approach in [14] to compute $CVaR_\alpha^{emp}(\bar{f})$

III. PRELIMINARIES ON KERNELS AND REPRODUCING KERNEL HILBERT SPACE (RKHS)

This section introduces the concept of a kernel, feature map and Reproducing Kernel Hilbert space (RKHS) which forms the backbone of our approach MMD-OPT. For a detailed exposition of these topics, refer [20].

A. Feature Maps and Kernels

Let \mathcal{X} be a measurable input space. Consider a vector $\mathbf{z} \in \mathcal{X}$ and a non-linear transformation ϕ called feature-map. We can use ϕ to embed \mathbf{z} in the RKHS \mathcal{H} ([20] Section II-B) as $\phi(\mathbf{z})$. A positive definite real-valued kernel function $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is related to the feature map through the so-called kernel trick $K(\mathbf{z}, \mathbf{z}') = \langle \phi(\mathbf{z}), \phi(\mathbf{z}') \rangle_{\mathcal{H}}$ ([20] Definition 2.1), where $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ is an inner product in \mathcal{H} . Intuitively, the inner product provides the distance between \mathbf{z}, \mathbf{z}' in RKHS. The kernel trick has the following implication: to compute the inner product $\langle \phi(\mathbf{z}), \phi(\mathbf{z}') \rangle_{\mathcal{H}}$, we do not have to explicitly compute $\phi(\mathbf{z})$. Instead, we just need to evaluate the kernel function over the pair \mathbf{z}, \mathbf{z}' as $K(\mathbf{z}, \mathbf{z}')$ and this will give us the required inner product. We use the Laplace Kernel throughout our implementation. It is given by:

$$K(\mathbf{z}, \mathbf{z}') = \exp\left(-\frac{\|\mathbf{z} - \mathbf{z}'\|_1}{\sigma}\right) \quad (13)$$

where $\|\cdot\|_1$ denotes the l^1 norm (also known as the *Manhattan* or *Taxicab* norm). Here $\sigma > 0$ is a tuneable bandwidth parameter.

B. From Data Points to Probability Distributions

We can generalize the concept of a feature map, kernel function and RKHS embedding to work on probability distributions as well (see Figure 4). Let u be a random variable with probability distribution p_u . The RKHS embedding of u (or p_u) is given by the following.

$$\mu[u] = E_{u \sim p_u}[\phi(u)] = E_{u \sim p_u}[K(u, \cdot)] \quad (14a)$$

$$= \int K(u, \cdot) dp_u(u) \quad (14b)$$

Essentially, (14a)-(14b) map probability distribution to a function in RKHS. For arbitrary distribution p_u , it is often

Algorithm 1 Reduced-Set Selection and Weight and Kernel Parameter Optimization using Cross-Entropy Method (CEM)

```

1  $e_{cem} =$  Maximum number of iterations,  $N =$  size of  $\mathbf{O}$ ,  $N' =$  size of reduced-set
2 Initialize mean and covariance  $\mathbf{v}_e, \Sigma_e$ , respectively, at  $e = 0$  for sampling and sorting  $\lambda$  and the kernel  $\sigma$  parameter
   respectively.
3 for  $e = 1, e \leq e_{cem}, e++$  do
4   Draw  $n_{cem}$  samples  $({}^1\lambda, {}^2\lambda, {}^m\lambda, \dots, {}^{n_{cem}}\lambda)$  and  $({}^1\sigma, {}^2\sigma, {}^m\sigma, \dots, {}^{n_{cem}}\sigma)$  from  $\mathcal{N}(\mathbf{v}_e, \Sigma_e)$ .
5    ${}^m\lambda = ({}^{m,1}\lambda, {}^{m,2}\lambda, \dots, {}^{m,N}\lambda)$  // The dimension of  ${}^m\lambda$  is equal to the rows of  $\mathbf{O}$ 
6   Initialize  $CostList = []$ 
7   for  $m = 1, m \leq n_{cem}, m++$  do
8      $({}^1\tau', {}^2\tau', {}^l\tau', \dots, {}^{N'}\tau') = F_{m\lambda}(\mathbf{O})$  // Reduced-Set selection formulated as sorting of the
       elements of  ${}^m\lambda$ 
9      ${}^l\beta^* \leftarrow$  Solve optimization (27c) with the reduced-set formed above and with kernel parameter  ${}^m\sigma$ 
10     $cost_{mmd} \leftarrow$  Compute MMD cost in (27c) using  ${}^l\beta^*$ 
11    append  $cost_{mmd}$  to  $CostList$ 
12  end
13   $EliteSet \leftarrow$  Select top  $n_{elite}$  samples of  $({}^m\lambda)_{m=1}^{m=n_{cem}}$  and  $({}^m\sigma)_{m=1}^{m=n_{cem}}$  with lowest cost from  $CostList$ .
14   $\mathbf{v}_{e+1}, \Sigma_{e+1} \leftarrow$  Update distribution based on  $EliteSet$  // Empirical mean and covariance estimate
     from the  $EliteSet$  samples
15 end
16 return  $({}^1\tau', {}^2\tau', \dots, {}^{N'}\tau')$ , weights  ${}^l\beta^*$  and optimal kernel hyperparameter  $\sigma$  corresponding to lowest cost in the
      $EliteSet$ 

```

intractable to compute the r.h.s of (14b). Thus, instead, we can resort to the empirical estimate of the RKHS embedding. Let ${}^j u$ be samples of u drawn from p_u . The empirical estimate of $\mu[u]$ is given by the following.

$$\hat{\mu}[u] := \sum_{j=1}^N \frac{1}{N} \phi({}^j u) = \sum_{j=1}^N \frac{1}{N} K({}^j u, \cdot) \quad (15)$$

1) *Functions of Random Variables:* We can extend the concept of RKHS embedding to functions of random variables as well. Let $g(u) \in \mathcal{X} \rightarrow \mathcal{Z}$ be an arbitrary function of a random variable u with p_g denoting the probability distribution of $g(u)$. Its RKHS embedding $\mu[g]$ and its empirical estimate $\hat{\mu}[g]$ is given by

$$\mu[g] = E_{u \sim p_u}[\phi(g(u))] = E_{u \sim p_u}[K(g(u), \cdot)] \quad (16a)$$

$$= \int K(g(u), \cdot) dp_u(u), \quad (16b)$$

$$\hat{\mu}[g(u)] := \sum_{j=1}^N \frac{1}{N} \phi(g({}^j u)) = \sum_{j=1}^N \frac{1}{N} K(g({}^j u), \cdot) \quad (16c)$$

Note that K in the above expressions (16a)-(16c) is the same as that used in (15).

One of the attractive features of RKHS embedding is that when constructed with characteristic Kernels (Gaussian, Laplacian, etc.), it can capture information about mean and moments of u and $g(u)$ up to infinite order. Furthermore, $\hat{\mu}[u], \hat{\mu}[g]$ are consistent estimators. That is they converge to the true RKHS embedding as we increase the sample size N . The RKHS embeddings $\hat{\mu}[u], \hat{\mu}[g]$ are also tightly coupled through the following theorem from [11].

Theorem 1: Suppose g is a continuous function and the kernel functions are continuous and bounded. Then, the following holds

$$\text{if } \hat{\mu}[u] \rightarrow \mu[u], \text{ then } \hat{\mu}[g] \rightarrow \mu[g] \quad (17)$$

The proof can be found in the Appendix of [11].

C. Maximum Mean Discrepancy and Improving Sample Efficiency

One of the core applications of RKHS embedding lies in quantifying similarity/difference between two probability distributions through a concept called Maximum Mean Discrepancy (MMD) defined in the following manner.

Definition 1: Let u and \tilde{u} be two random variables with probability distributions p_u and $p_{\tilde{u}}$ respectively. The MMD between the two distributions is given by

$$MMD(p_u, p_{\tilde{u}}) = \|\mu[u] - \mu[\tilde{u}]\|_{\mathcal{H}}^2 = \langle \mu[u] - \mu[\tilde{u}], \mu[u] - \mu[\tilde{u}] \rangle \quad (18)$$

where $\|(\cdot)\|_{\mathcal{H}}$ is the RKHS norm and can be expressed through inner products, which in turn can be evaluated using the Kernel trick. The distribution p_u and $p_{\tilde{u}}$ are same if and only if $\|\mu[u] - \mu[\tilde{u}]\|_{\mathcal{H}}^2 = 0$ [11], [21]. Typically, it is intractable to obtain the exact MMD and thus, we resort to its empirical estimate given by the following

$$MMD^{emp}(p_u, p_{\tilde{u}}) = \|\hat{\mu}[u] - \hat{\mu}[\tilde{u}]\|_{\mathcal{H}}^2 \quad (19)$$

IV. MAIN ALGORITHMIC RESULTS

This section presents our main algorithmic contributions. An overview of our approach is shown in Fig. 5. The cornerstone of our work is the *reduced-set selection* (Alg. 1) which allows us to use only a subset of predicted obstacle trajectories to compute our MMD-based collision risk surrogate. This in turn, drastically reduces the number of collision-checks that we need to perform. The optimal *reduced-set* is then passed to the downstream sampling-based optimizer (Alg. 2) wherein we sample ego-vehicle trajectories in the Frenet-frame from a Gaussian distribution and evaluate their associated costs (including collision risk). The sampling distribution is updated in an iterative fashion to drive the costs down at each iteration.

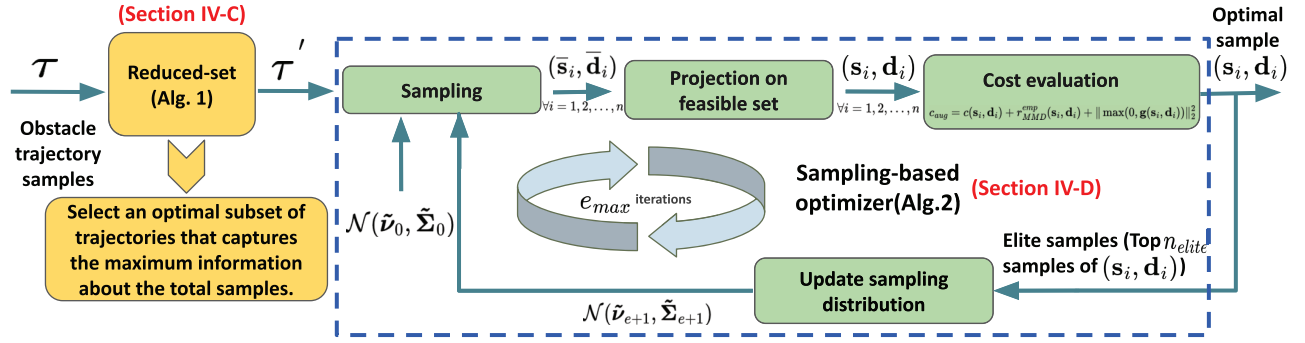


Fig. 5. Overview of MMD-OPT. It consists of two main parts. The first part is the so-called *reduced-set* algorithm that tells us which obstacle trajectories can be discarded without compromising the ability of our MMD-based surrogate to capture the collision risk in a given scene. The second part is a sampling-based optimizer that minimizes the MMD-based collision risk cost. Our sampling based optimizer also consists of novel components that exploits the problem structure for computational efficiency.

Algorithm 2 Sampling-Based Optimizer to Solve (5)–(7) With r_{MMD}^{emp} as the Risk Cost

```

1  $e_{max}$  = Maximum number of iterations, reduced-set
   samples of obstacle trajectories ( ${}^1\tau', {}^2\tau', \dots, {}^N\tau'$ ) and
   kernel parameter  $\sigma$  from Alg.1
2 Initiate mean  $\tilde{\mathbf{v}}_e, \tilde{\Sigma}_e$ , at  $e = 0$  for sampling
   Frenet-Frame behavioural inputs  $\mathbf{b}$ 
3 for  $e = 1, e \leq e_{max}, e++$  do
4   Draw  $n$  samples ( $\mathbf{b}_1, \mathbf{b}_2, \dots, \mathbf{b}_n$ ) from
      $\mathcal{N}(\tilde{\mathbf{v}}_e, \tilde{\Sigma}_e)$ 
5   Initialize  $CostList = []$ 
6   Query Frenet-planner for  $\forall \mathbf{b}_i$ :
      $(\bar{\mathbf{s}}_i, \bar{\mathbf{d}}_i) = \text{Frenet Planner}(\mathbf{b}_i)$ 
7   Project to Constrained Set
8
9      $(\mathbf{s}_i, \mathbf{d}_i) = \arg \min_{\mathbf{s}_i, \mathbf{d}_i} \frac{1}{2} \|\bar{\mathbf{s}}_i - \mathbf{s}_i\|_2^2$ 
        $+\frac{1}{2} \|\bar{\mathbf{d}}_i - \mathbf{d}_i\|_2^2$ 
        $\mathbf{h}(\mathbf{s}_i^{(q)}, \mathbf{d}_i^{(q)}) = \mathbf{0}$ 
        $\mathbf{g}(\mathbf{s}_i^{(q)}, \mathbf{d}_i^{(q)}) \leq 0$ 
10  Define constraint residuals:  $\|\max(0, \mathbf{g}(\mathbf{s}_i, \mathbf{d}_i))\|_2^2$ 
      $ConstraintEliteSet \leftarrow$  Select top  $n_c$  samples of
        $\mathbf{b}_i, (\mathbf{s}_i, \mathbf{d}_i)$  with lowest constraint residual norm.
11  Define  $c_{aug} =$ 
      $c(\mathbf{s}_i, \mathbf{d}_i) + r_{MMD}^{emp}(\mathbf{s}_i, \mathbf{d}_i) + \|\max(0, \mathbf{g}(\mathbf{s}_i, \mathbf{d}_i))\|_2^2$ 
      $cost \leftarrow c_{aug}$ , over  $ConstraintEliteSet, \forall i$ 
12  append each computed  $cost$  to  $CostList$ 
13   $EliteSet \leftarrow$  Select top  $n_{elite}$  samples of
      $(\mathbf{b}_i, \mathbf{s}_i, \mathbf{d}_i)$  with lowest cost from  $CostList$ .
14   $(\tilde{\mathbf{v}}_{e+1}, \tilde{\Sigma}_{e+1}) \leftarrow$  Update distribution based on
      $EliteSet$ 
15
16 end
17
18 return Frenet parameter  $\mathbf{b}_i$  and  $(\mathbf{s}_i, \mathbf{d}_i)$ 
   corresponding to lowest cost in the  $EliteSet$ 

```

In the next few sub-sections, we present the different components from Fig. 5. We begin by presenting our surrogate for collision risk.

A. Collision Risk as Difference of Two Distributions

The distribution of f and \bar{f} , denoted by p_f and $p_{\bar{f}}$ respectively, cannot be characterized analytically, even if we assume that the predicted trajectory of the neighbouring vehicle is drawn from a Gaussian distribution. In general, p_τ can be arbitrarily complex, making it even more challenging to make any assumption on the nature of p_f and $p_{\bar{f}}$. However, we can intuitively reason about some desired characteristics that the latter two distribution should have.

- As shown in Fig. 3, we want to shift the mass of distribution p_f to the left of $f = 0$. The higher the mass is to the left of $f = 0$, the lower the risk $r = P(f \geq 0)$.
- As the mass of p_f gets shifted to the left of $f = 0$, $p_{\bar{f}}$ converges to a Dirac-Delta distribution p_δ . This follows directly from the definition of \bar{f} in (10).
- We can re-imagine risk-aware trajectory optimization as the problem of finding the optimal ego vehicle trajectory (\mathbf{s}, \mathbf{d}) that re-shapes $p_{\bar{f}}$ to look similar to a Dirac-Delta distribution. This, in turn, would ensure that the mass of p_f gets shifted to the left of $f = 0$.

Based on the above insights, we propose the following novel surrogate for collision risk

$$r_{MMD} = MMD(p_{\bar{f}}, p_\delta) = \|\mu[\bar{f}] - \mu[\delta]\|_{\mathcal{H}}^2 \quad (20)$$

where $\mu[\bar{f}]$ is the RKHS embedding of \bar{f} (or $p_{\bar{f}}$) defined in the following manner

$$\mu[\bar{f}] = \int K(\bar{f}, \cdot) dp_{\bar{f}}(\bar{f}), \quad \hat{\mu}[\bar{f}] = \sum_{j=1}^N \frac{1}{N} \phi(j\bar{f}) \quad (21)$$

It is connected to the RKHS embedding of obstacle trajectory distribution p_τ defined as:

$$\mu[\tau] = \int K(\tau, \cdot) dp_\tau(\tau), \quad \hat{\mu}[\tau] = \sum_{j=1}^N \frac{1}{N} \phi(j\tau) \quad (22)$$

Similarly, $\mu[\delta]$ is the RKHS embedding of a random variable $\delta \sim p_\delta$. As r_{MMD} decreases, $p_{\bar{f}}$ becomes similar to p_δ . In the limiting case, as $r_{MMD} \rightarrow 0$, the distribution $p_{\bar{f}} \rightarrow p_\delta$ and consequently $P(f \geq 0) \rightarrow 0$.

Since the true r_{MMD} is intractable to compute for arbitrary p_τ , we will be working with the following empirical estimate-computed through sample-level information

$$r_{MMD}^{emp} = \|\hat{\mu}[\bar{f}] - \hat{\mu}[\delta]\|_{\mathcal{H}}^2 \quad (23)$$

Remark 5: By Theorem 1, $\hat{\mu}[\tau] \rightarrow \mu[\tau]$ implies $\hat{\mu}[\bar{f}] \rightarrow \mu[\bar{f}]$ and consequently $r_{MMD}^{emp} \rightarrow r_{MMD}$

Remark 6: We can approximate p_δ through a Gaussian $\mathcal{N}(0, \epsilon)$, with an extremely small covariance $\epsilon (\approx 10^{-5})$.

1) *Finite Sample Guarantees:* A key advantage of r_{MMD}^{emp} is that it inherits the excellent convergence and finite sample guarantees associated with RKHS embedding and MMD. Specifically, how far r_{MMD}^{emp} is from the true r_{MMD} can be quantified by the following Theorem that is a direct application of the result presented in [12].

Theorem 2: Let \bar{f}, δ be the collision constraint violation and Dirac-Delta random variables respectively. Let $p_{\bar{f}}, p_\delta$ be the respective probability distributions of \bar{f} and δ . Given samples $X := \{\bar{f}_j\}_{j=1}^N$ and $Y := \{\delta_j\}_{j=1}^N$ i.i.d from $p_{\bar{f}}$ and p_δ , respectively, and assuming $0 \leq K(\xi_1, \xi_2) \leq c_o, \forall \xi_1, \xi_2$ and some constant c_o , then:

$$P_{X,Y} \left\{ |r_{MMD}^{emp} - r_{MMD}| \leq 2 \left(2(c_o/N)^{\frac{1}{2}} + \epsilon \right) \right\} > 1 - \exp\left(\frac{-\epsilon^2 N}{4c_o}\right) \quad (24)$$

where $P_{X,Y}$ denotes the probability over the N -sample X and Y , and $\epsilon > 0$ is an arbitrary parameter controlling the probability bound. $K(x, y)$ denotes the kernel function.

Theorem 2 ensures that as the sample size N increases, the probability that the absolute difference between r_{MMD}^{emp} and r_{MMD} tends to a small value approaches 1. In other words, with increase in samples, not only the estimate r_{MMD}^{emp} improves, the stochasticity or the variance associated with it also reduces. In Section VI, we show Theorem 2 directly translates to lower worst-case collision-rate with r_{MMD}^{emp} as compared to r_{CVaR}^{emp} and r_{SAA} .

B. Advantages of r_{MMD}^{emp} as a Collision Risk Surrogate

Just like r_{SAA} (Eqn. (11)) and r_{CVaR} (Eqn. (12)), our r_{MMD}^{emp} also requires evaluating the collision constraint residual function \bar{f} over predicted obstacle trajectories. This can be computationally intensive if the collision checks are expensive. The key feature that differentiates r_{MMD}^{emp} from r_{SAA} and r_{CVaR}^{emp} is that we can use the tools from RKHS embedding and MMD to systematically improve its sample complexity. In other words, we can use only a subset $N' \ll N$ of the predicted obstacle trajectories to evaluate \bar{f} and yet incur only minimum degradation in estimation of r_{MMD}^{emp} . We call this subset as *reduced-set* following the notation from [11] and [12].

Let ${}^l\tau', l = 1, 2, 3, \dots, N'$ be the *reduced-set* formed by sub-selecting N' samples out of N samples of τ . Moreover, let us attach weights ${}^l\beta$ with every ${}^l\tau'$ such that $\sum_l {}^l\beta = 1$. Let ${}^l\bar{f}'$ be the samples of \bar{f} evaluated on ${}^l\tau'$. The empirical RKHS embeddings of p_τ and $p_{\bar{f}}$ computed using ${}^l\tau'$ and ${}^l\bar{f}'$ are given by:

$$\hat{\mu}[\tau'] = \sum_l {}^l\beta\phi({}^l\tau'), \quad \hat{\mu}[\bar{f}'] = \sum_l {}^l\beta\phi({}^l\bar{f}') \quad (25)$$

Now, [11] showed that Theorem 1 can also be leveraged to ensure that if $\hat{\mu}[\tau']$ is close to $\hat{\mu}[\tau]$ in MMD metric, then $\hat{\mu}[\bar{f}']$ will be close to $\hat{\mu}[\bar{f}]$ in the same metric. Consequently, r_{MMD}^{emp} computed using $\hat{\mu}[\bar{f}']$ which requires N' predicted obstacle trajectories will be closer to that computed using $\hat{\mu}[\bar{f}]$ on the larger set of N samples. Alternately, if we reduce the sample size in a way that we minimize $\|\hat{\mu}[\tau'] - \hat{\mu}[\tau]\|_{\mathcal{H}}^2$, the loss of accuracy in estimation of r_{MMD}^{emp} will be minimal. In the next sub-section, we present several ways by which this can be achieved.

C. Reduced-Set

There are three ways to minimize $\|\hat{\mu}[\tau'] - \hat{\mu}[\tau]\|_{\mathcal{H}}^2$. First, we can choose the *reduced-set* in a careful manner. Second, we optimize the values of ${}^l\beta$. Finally, we can also choose the right set of kernel parameters σ . In the subsequent subsections, we first present a simple method that adopts only the second approach and is a straightforward application of [11]. Subsequently, we present our key result that adopts all three ways to minimize $\|\hat{\mu}[\tau'] - \hat{\mu}[\tau]\|_{\mathcal{H}}^2$.

1) *Fixed Reduced-Set, Optimal Weights:* In this approach, we assume that the N' *reduced-set* samples ${}^l\tau'$ are simply a random sub-selection from N samples of τ . We then optimize the weights of the ${}^l\tau'$ through the following optimization problem.

$$\min_{{}^l\beta} \left\| \frac{1}{N} \sum_{i=1}^{i=N} \phi({}^i\tau) - \sum_{l=1}^{l=N'} {}^l\beta\phi({}^l\tau') \right\|_{\mathcal{H}}^2 \quad (26a)$$

$$\sum_l {}^l\beta = 1 \quad (26b)$$

Using the kernel-trick, optimization (26a)-(26b) can be rephrased as an equality-constrained QP. Essentially, (26a)-(26b) can be seen as distilling the information from a significantly larger sample set of predicted obstacle trajectories to the *reduced-set* by re-weighting the importance of the samples of the latter.

2) *Optimal Reduced-Sets, Weights and Kernel Parameters:* In this approach, we simultaneously optimize the selection of the *reduced-set*, the weights of its individual samples and kernel parameters. Our approach takes the form of a bi-level optimization (27a)-(27c), wherein \mathbf{O} is a matrix with N rows, each containing a sample of τ drawn from p_τ .

$$\min_{\lambda, \sigma} \left\| \frac{1}{N} \sum_{i=1}^{i=N} \phi({}^i\tau) - \sum_{l=1}^{l=N'} {}^l\beta^* \phi({}^l\tau') \right\|_{\mathcal{H}}^2 \quad (27a)$$

$$F_\lambda(\mathbf{O}) = ({}^1\tau', {}^2\tau', \dots, {}^{N'}\tau') \quad (27b)$$

$${}^l\beta^* = \arg \min_{{}^l\beta} \left\| \frac{1}{N} \sum_{i=1}^{i=N} \phi({}^i\tau) - \sum_{l=1}^{l=N'} {}^l\beta\phi({}^l\tau') \right\|_{\mathcal{H}}^2$$

$$\text{s.t. } \sum_l {}^l\beta = 1 \quad (27c)$$

In (27a)-(27c), F_λ is a selection function that chooses N' rows out of \mathbf{O} to form the *reduced-set*. It is parameterized by λ . That is, different λ leads to different *reduced-set*. We discuss the algebraic form for F_λ later in this section. But, it is worth pointing out that if we fix λ and consequently F_λ (e.g




Fig. 6. Validation of performance of Alg. 1. Fig. (a) shows the progression of outer-level cost (27a) across iterations. Fig. (b) validates the efficacy of our optimal *reduced-set* selection. We sampled different random *reduced-sets* and computed the average, minimum and maximum of MMD $\|\hat{\mu}[\tau'] - \hat{\mu}[\tau]\|_{\mathcal{H}}^2$. As shown in Fig. (b), our optimal *reduced-set* leads to an $\|\hat{\mu}[\tau'] - \hat{\mu}[\tau]\|_{\mathcal{H}}^2$ that is very close to what can be achieved by exhaustively trying out different randomly selected *reduced-set*. Fig. (b) validates our hypothesis based on Theorem 1 and insights from [11]. As $\hat{\mu}[\tau'] \rightarrow \hat{\mu}[\tau]$, the MMD $\|\hat{\mu}[\tau'] - \hat{\mu}[\tau]\|_{\mathcal{H}}^2 \rightarrow 0$. We can once again observe that our optimal *reduced-set* performs close to the best-case performance than can be hoped to be achieved by trying-out different random *reduced-sets*. Fig. (c) verifies that our optimal *reduced-set* selection ensures minimal loss in estimation of r_{MMD}^{emp} while reducing the sample size. To this end, we computed $\Delta_{r_{MMD}^{emp}}$ through (29) which captures the estimation error between r_{MMD}^{emp} computed over a large set and the smaller *reduced-set*. As can be seen, r_{MMD}^{emp} computed over the optimal *reduced-set* samples perform close to the average performance that can be achieved by exhaustively trying out different random choices of the *reduced-set*. It should be noted that in real-world setting, exhaustive sampling is not possible and our bi-level optimization (27a)-(27c) provides a one-shot solution that is as good as exhaustive sampling in expectation.

just a random λ) and the kernel parameter σ , then we recover the earlier approach described in (26a)-(26b).

As can be seen, the inner optimization (27c) is defined over just the weights β for a fixed λ and kernel parameter σ . The outer optimization in turn optimizes in the space of λ and kernel parameter in order to reduce the MMD cost associated with optimal β^* . Our bi-level approach provides two pronged benefits. First, it gets rid of the variance associated with randomly selecting the *reduced-set*. Second, our approach leads to $\hat{\mu}[\tau']$ that has lower MMD difference from $\hat{\mu}[\tau]$ as compared to that achieved from the previous method of Section IV-C1. This in turn, improves the embedding of $\hat{\mu}[f']$ and consequently the estimate of r_{MMD}^{emp} .

a) *Selection Function* F_λ : Let $\lambda \in \mathbb{R}^N$ be an arbitrary vector and λ_t be its t^{th} element. Assume that $|\lambda_t|$ encodes the value of choosing the t^{th} row of \mathbf{O} to form the *reduced-set*. That is, the larger the $|\lambda_t|$, the higher the impact of choosing the t^{th} row of \mathbf{O} in minimizing (27a). With these constructions in place, we can now define F_λ through the following sequence of mathematical operations.

$$S = \text{Argsort}\{|\lambda_1|, |\lambda_2|, \dots, |\lambda_N|\} \quad (28a)$$

$$\mathbf{O}' = [\mathbf{O}_{t_1}, \mathbf{O}_{t_2}, \dots, \mathbf{O}_{t_N}], \forall t_e \in S \quad (28b)$$

$$F_\lambda(\mathbf{O}) = \mathbf{O}'_{N-N':N} = ({}^1\tau', {}^2\tau', \dots, {}^{N'}\tau') \quad (28c)$$

As can be seen, we first sort in increasing value of $|\lambda_t|$ and compute the required indices t_e . We then use the same indices

to shuffle the rows of \mathbf{O} in (28b) to form an intermediate variable \mathbf{O}' . Finally, we choose the last N' rows of \mathbf{O}' as our *reduced-set* in (28c). Intuitively, (28a)-(28c) parameterizes the sub-selection of the rows of \mathbf{O} to form the *reduced-set*. That is, different λ gives us different *reduced-sets*. Thus, a large part of solving (27a)-(27c) boils down to arriving at the right λ . We achieve this through a combination of gradient-free search and quadratic programming (QP). We present the details of our approach next.

b) *Solution Process*: Our approach for solving (27a)-(27c) is presented in Alg. 1 that combines gradient-free Cross-Entropy Method (CEM) [22] and QP. The input to the algorithm is the set \mathbf{O} and the required size of the reduced-set N' . The output is the *reduced-set* samples, their weights and the optimal kernel parameter σ . Alg. 1 begins by initializing a Gaussian distribution (Line 2) with mean v_e and covariance Σ_e at iteration $e = 0$. The loop defined between lines 3-15 is the outer CEM optimizer. At each iteration of CEM, we sample n_{cem} batch of λ and σ . Note that the dimension of λ is equal to the number of obstacle trajectory samples contained in \mathbf{O} . Subsequently, between Lines 7-12, we loop through each sampled λ and σ and use them to formulate and solve the inner optimization problem. First, we use $F_\lambda(\mathbf{O})$ as defined in (28a)-(28c) to obtain the *reduced-set* samples. These are then used to solve optimization (27c) along with σ as the kernel parameter to obtain β^* . At line 10, we compute

the optimal MMD resulting from our choice of $l\tau$ samples, $l\beta^*$ and $m\sigma$ and append it to our cost-list. Lines 13-14 refines the sampling distribution of the outer loop. We select the top n_{elite} samples of $m\lambda$ and $m\sigma$ that led to the lowest MMD cost and call it the *EliteSet*. Line 14 estimates the empirical mean and covariance over the $m\lambda$ and $m\sigma$ contained in the *EliteSet*. Alg. 1 uses GPU parallelization to achieve computational efficiency. Specifically, the inner QP loop in lines 7-12 can be performed in parallel over different m . Moreover, we also leverage the fact that the equality constrained QP (27c) have a closed-form solution. Fig. 6 validates the efficacy of Alg. 1 in solving (27a)-(27c): the outer-level cost gradually decreases with iterations.

c) Validating the Efficacy of the Optimal Reduced-Set:

To empirically showcase the importance of our optimal *reduced-set*, we sampled 500 obstacle trajectories to form \mathbf{O} . We then sub-selected N' rows from it, either randomly or by solving the bi-level optimization (27a)-(27c). We perform a large number of random sub-selections to compute the average, minimum and maximum of $\|\hat{\mu}[\tau'] - \hat{\mu}[\tau]\|_{\mathcal{H}}^2$. Moreover, we compute the same MMD over the optimal *reduced-set* as well.

As can be seen from Fig. 6 (b), $\|\hat{\mu}[\tau'] - \hat{\mu}[\tau]\|_{\mathcal{H}}^2$ over the optimal *reduced-set* performs close to the best-case performance that we can obtain by exhaustively trying out different random *reduced-sets*. The performance difference is more stark at lower N' . It should be noted that in real-world setting, exhaustive search is prohibitive and thus, (27a)-(27c) provides a computationally efficient one-shot solution. Fig. 6 (c) shows that the advantages of optimal *reduced-set* translates to empirical embedding $\hat{\mu}[\bar{f}']$ as well, with $\|\hat{\mu}[\bar{f}'] - \hat{\mu}[\bar{f}]\|_{\mathcal{H}}^2$ being only slightly inferior to the best-case value obtained with exhaustive sampling of *reduced-set*. Fig. 6 validates our hypothesis that minimizing $\|\hat{\mu}[\tau'] - \hat{\mu}[\tau]\|_{\mathcal{H}}^2$ and $\|\hat{\mu}[\bar{f}'] - \hat{\mu}[\bar{f}]\|_{\mathcal{H}}^2$ ensures minimal loss in the estimation of r_{MMD}^{emp} . To this end, we computed

$$\Delta_{r_{MMD}} = \left\| \left\| \sum_{j=1}^N \frac{1}{N} \phi(\bar{f}^j) - \sum_{j=1}^N \frac{1}{N} \phi(\delta^j) \right\|_{\mathcal{H}}^2 - \left\| \sum_{l=1}^{N'} \beta^l \phi(\bar{f}^l) - \sum_{l=1}^{N'} \frac{1}{N'} \phi(\delta^l) \right\|_{\mathcal{H}}^2 \right\| \quad (29)$$

The first term is the r_{MMD}^{emp} computed on the larger set of N samples, while the second term is the same computed on the smaller *reduced-set*. Like before, we constructed the *reduced-set* either randomly or by solving (27a)-(27c). As shown in Fig. 6, the $\Delta_{r_{MMD}}$ obtained with the optimal *reduced-set* is close to the average error that can be achieved by exhaustively trying different choices of the *reduced-set*.

d) *Physical Manifestation of Optimal Reduced-Set:* We now present additional empirical analysis which sheds more light into the concept of optimal *reduced-set*. Specifically, we analyze if the obstacle trajectory samples contained in the optimal *reduced-set* have a physical significance. To this end, consider Fig. 7 which shows a synthetic multi-modal obstacle trajectory distribution that spans two different lane-change intents. This distribution was generated by sampling set-points for lateral offsets (from the lane center-line) and

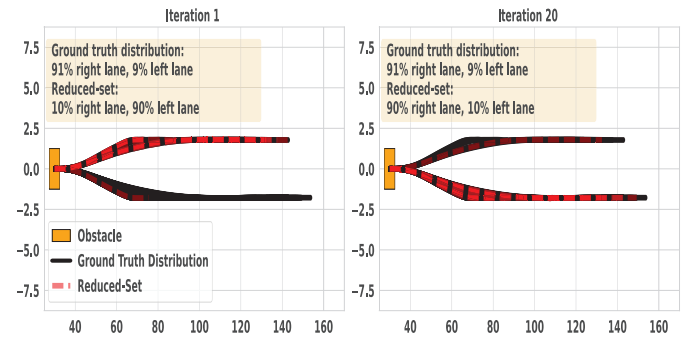


Fig. 7. The figures demonstrate the physical significance of the trajectory samples used to form the optimal *reduced-set*. We generate a synthetic multi-modal trajectory distribution that captures the different lane-change intents of a dynamic obstacle (orange rectangle). More importantly, due to our specific generation procedure, we know the approximate probability of each lane-change intent. When Alg. 1 starts, the optimal *reduced-set* in the first iteration consists mostly of the less likely samples due to the random initialization of the Gaussian distribution from which λ is sampled. **For example in the left figure, the *reduced-set* samples are mostly formed by the left lane-change maneuvers which is less likely.** But as Alg. 1 makes progress, the composition significantly changes. For example, in the right figure, the majority of the *reduced-set* samples switch from left-lane to the right-lane maneuvers. Thus, by just minimizing MMD between two sets of trajectory samples, Alg. 1 is able to implicitly infer which trajectory samples are more likely and accordingly split the composition of the optimal *reduced-set* between the less and more likely samples.

forward velocity from a binomial distribution and a Gaussian Mixture Model respectively and mapping those to trajectory distribution using the planner presented in Appendix A. Due to discrete nature of the lateral offset set-point distribution, we can approximately infer which of the lane-change maneuvers of Fig. 7 are more likely. We now apply Alg. 1 to choose a subset of trajectory samples to form the optimal *reduced-set*. We visualize the chosen samples in Fig. 7. As can be seen, at the first iteration of Alg. 1, the optimal *reduced-set* is primarily formed by less likely samples from the distribution. But as Alg. 1 progresses and the outer MMD cost (27a) goes down, the composition of the optimal *reduced-set* drastically changes. Specifically, more samples of more likely trajectories are found in the optimal *reduced-set* at later iterations. In fact, the split into more and less likely trajectory samples in the optimal *reduced-set* closely follow the ground-truth probability. We observed this trend consistently across all our simulations, and we believe, it points to an interesting emergent behavior in Alg. 1. To summarize, although, Alg. 1 does not use the actual probability of a trajectory sample in any way, it finds an implicit way of inferring the more likely samples by just minimizing the MMD between two sets of trajectory samples. This in turn, ensures that r_{MMD}^{emp} constructed on the optimal *reduced-set* demonstrates excellent ability to capture collision risk and thus lead to safer motions.

D. Trajectory Optimizer

We minimize (5) subject to (6)-(7) with r_{MMD}^{emp} as our risk cost, through sampling-based optimization that combines features from CEM [22], Model Predictive Path Integral (MPPI) [23] and its modern variants [24]. This class of optimizers operates by drawing at each iteration, a batch of ego-vehicle

trajectories from a parametric distribution and evaluating the cost over all of them. Subsequently, the cost values are used to adapt the parameters of the sampling distribution for the next iteration. Typically, these optimizer are used in unconstrained setting. Thus, to account for inequality constraints (7), we embed a projection optimizer into the sampling pipeline that is responsible for pushing the sampled trajectories towards feasible region (e.g lane boundary) before evaluating their cost.

Our proposed sampling-based optimizer is presented in Alg. 2.¹ The input to our algorithm is the allowable iteration limit and the optimal *reduced-set* samples computed from Alg. 1. We leverage the problem structure by sampling trajectories in Frenet-Frame. To this end, we first sample n behavioural inputs \mathbf{b}_i such as desired lateral offsets and longitudinal velocity set-points from a Gaussian distribution (line 4). These are then fed into a Frenet space planner inspired by [25] in line 6, effectively mapping the distribution over behavioural inputs to that over trajectories. We present the details about the Frenet planner in Appendix. The obtained trajectories are then passed in line 7 to a projection optimization that pushes the trajectories towards a feasible region. Our projection optimizer is built on top of [26], [27] and can be easily parallelized over GPUs. In line 8, we evaluate the constraint residuals over the projection outputs ($\mathbf{s}_i, \mathbf{d}_i$). In line 9, we select top $n_c (< n)$ projection outputs with the least constraint residuals. We call this set the *ConstraintEliteSet*. In line 10, we form c_{aug} that consists of the primary cost (5) augmented with the constraint residual. In line 11-12, we append c_{aug} to the *CostList*. In line 13, we extract the top $n_{elite} (\leq n_c)$ trajectory samples form *ConstraintEliteSet* that led to the lowest c_{aug} . We call this the *EliteSet*. In line 14, we update the sampling distribution of behavioural inputs based on the cost samples collected in the *EliteSet*. The exact formula for mean and covariance update is presented in (30a)-(30b). Herein, the constants γ and η are the temperature and learning rate respectively. The notion of *EliteSet* is brought from classic CEM, while the mean and covariance update follow from MPPI [23], and [24]. An important feature of Alg. 2 is we have effectively encoded long horizon trajectories with a low dimensional behavioural input vector samples \mathbf{b}_i . This, in turn, improves the computational efficiency of our approach.

$$\tilde{\mathbf{v}}_{e+1} = (1 - \eta)\tilde{\mathbf{v}}_e + \eta \frac{\sum_{i=1}^{i=n_{elite}} \omega_i \mathbf{b}_i}{\sum_{i=1}^{i=n_{elite}} \omega_i}, \quad (30a)$$

$$\tilde{\Sigma}_{e+1} = (1 - \eta)\tilde{\Sigma}_e + \eta \frac{\sum_{i=1}^{i=n_{elite}} \omega_i (\mathbf{b}_i - \tilde{\mathbf{v}}_{e+1})(\mathbf{b}_i - \tilde{\mathbf{v}}_{e+1})^T}{\sum_{i=1}^{i=n_{elite}} \omega_i} \quad (30b)$$

$$\omega_i = \exp \frac{-1}{\gamma} (c_{aug}(\mathbf{s}_i, \mathbf{d}_i)) \quad (30c)$$

E. Extension to Multiple Obstacles

We follow a very basic approach for extending our planning pipeline to multiple obstacles. For each obstacle, we compute individual optimal *reduced-sets* and kernel parameters through

¹We use right subscript to denote ego trajectory samples generated at each iteration of Alg. 2 and left superscript to denote obstacle trajectories sampled from the trajectory predictor.

Alg. 1. Since each obstacle computation is decoupled from the others, this step can potentially be parallelized. In Alg. 2, we compute r_{MMD}^{emp} for each obstacle and add them to estimate the total collision risk.

V. CONNECTIONS TO EXISTING WORKS

In this section, we connect our approach to current state-of-the-art (SOTA), while also pointing out how we go beyond them. For the ease of exposition, we have classified the related works into specific categories.

A. Deterministic Approaches

Motion Planning is one of the most well-studied problems in autonomous driving. A variety of techniques based on potential-field based methods [28], spatio-temporal corridor generation [29] and trajectory optimization [27] have been proposed. A large section of the existing works including those cited here assume a deterministic environment, where the future trajectories of the obstacles/neighboring vehicles are precisely known. However, such assumptions do not hold in real-world setting as the intents of the other vehicles cannot be known a priori.

Improvement over SOTA: Our proposed approach provides a stronger safety guarantee in a more realistic setting where the trajectories of the obstacles have a general structure in the form of a probability distribution.

B. Chance Constrained Optimization

Instead of minimizing risk r (as defined in (9)), if we add constraints of the form $r \leq \zeta$ in (5)-(7), for some constant ζ , we obtain the setting of chance-constrained optimization (CCO). The core effort in CCO lies in computing tractable surrogate of the chance-constraints. To this end, $r_{SAA} \leq \zeta$, based on (11) is one such potential surrogate used in works like [13]. Similarly, a CVaR estimate defined over the collision constraint function f (recall eqn. 8) is used in [18] as a substitute for chance-constraints.

An alternative surrogate for chance-constraints namely the so-called scenario approximations is also often used [30], [31], [32], [33], [34]. Here $r \leq \zeta$ is replaced with deterministic scenario constraints of the form ${}^j f(\mathbf{s}, \mathbf{d}, {}^j \boldsymbol{\tau}) \geq 0$, defined with j^{th} sample of obstacle trajectory prediction. Although conceptually simple, scenario-approximation requires adding a large number of constraints to the optimization problem that can be computationally prohibitive. Moreover, as we accommodate more scenario constraints, the overall solution quality can drastically deteriorate [35].

Improvement over SOTA: Our approach can be seen as a relaxation of CCO where we minimize risk instead of constraining it to γ level. This alleviates the problem of infeasibility in case γ is set too low. Moreover, our approach puts no restriction in the underlying distribution of the uncertainty, which in our case stems from the obstacle trajectory prediction. In Section VI-G, we also show that our approach outperforms baseline scenario approach for CCO.

C. Risk-Aware Planning

Our work defines collision risk using the collision constraint residual function \bar{f} (recall (10)). A similar set-up can be found in [14], wherein the authors used the empirical CVaR estimate of \bar{f} as a risk estimate. Authors in [36] define risk in terms of spatio-temporal overlap of ego-vehicle and obstacle footprint and severity of collision. It is simple to observe that \bar{f} actually captures the footprint overlap in space and time. Several other risk metrics like time-to-collision [37] can also be used as a risk metric. However, estimating it for a given complex distribution of obstacle trajectories could be challenging.

Improvement over SOTA: One key differentiating feature between our approach and all the above cited works is the focus on sample complexity. In fact, we show in Section VI that our collision risk r_{MMD}^{emp} can ensure safety with much fewer number of samples of obstacle trajectories than existing risk metrics based on CVaR and SAA.

D. RKHS Embedding for Stochastic Trajectory Optimization

There have been prior works that have also tried to leverage some concepts from RKHS embedding and MMD for either chance-constrained or risk minimization. The authors in [35] use an MMD based rejection sampling to make the scenario approximations of chance-constraints less conservative. A recent work [38] uses the MMD to define ambiguity sets in distributionally robust chance-constrained problems.

Improvement over SOTA: Our work departs from [35] and [38] in a way that it constructs the risk surrogate itself on the top of MMD. This opened the doors to leveraging the power of reduced-set in a much more prominent manner than [35], [38]. Our current work extends [16], [39], [40], to simultaneously optimize the reduced-set selection, weights of the individual sample and the kernel parameters. Especially, in [39], the reduced-set selection and weights optimization were decoupled from each other with the kernel parameter being manually tuned to a fixed value. Our current approach not only leads to a better collision risk estimate but is also superior in terms of the final trajectory quality.

VI. VALIDATION AND BENCHMARKING

This section aims to answer the following research questions

- 1) **Q1:** How is the sample efficiency of r_{MMD}^{emp} as compared to CVaR and SAA estimate of collision constraint residual distribution? This is covered in Sections VI-C, VI-D, VI-E
- 2) **Q2:** What is the improvement in sample efficiency obtained by defining r_{MMD}^{emp} over an optimal *reduced-set* as compared to randomly selected one? Refer to Section VI-F
- 3) **Q3:** What is the scalability of r_{MMD}^{emp} based risk-aware trajectory optimization with respect to the number of obstacles (Section VI-J) ?

A. Implementation Details

We implemented our optimal *reduced-set* selection Alg. 1 and Alg. 2 in Python using JAX [41] library as our GPU-accelerated linear algebra back-end. Our collision constraint

function f_k is designed assuming that the ego-vehicle and the obstacle are modeled as axis-aligned ellipses with size (a_1, a_2) . Thus, it has the following mathematical form

$$f_k := -\frac{(s_k - s_{o,k})^2}{a_1^2} - \frac{(d_k - d_{o,k})^2}{a_2^2} + 1 \quad (31)$$

where $\tau_k = (s_{o,k}, d_{o,k})$ is the position of the obstacle at time-step k and is obtained from the trajectory predictor. Stacking τ_k at different time steps give us τ .

The cost term c in (5) has the following form

$$c(\mathbf{s}^{(q)}, \mathbf{d}^{(q)}) = c_\theta + c_v + c_a + c_{lane} \quad (32)$$

with

$$c_\theta = w_{\theta,1} \sum_k \max(0, |\theta_k| - \theta_{max}) + w_{\theta,2} \theta_k^2 + w_{\theta,3} \dot{\theta}_k^2 + w_{\theta,4} \ddot{\theta}_k^2 \quad (33)$$

The variable θ_k as defined earlier, is the steering angle. Using the motion-model (2)-(4), we can express θ_k as a function of the the first and second derivatives of the position level trajectory. The different $w_{\theta,i}$ are weights used to trade-off the effect of the cost terms with respect to each other. The derivatives of the θ_k are approximated through finite differencing.

The cost, c_v , c_a , c_{lane} are defined as

$$c_v = \sum_k (\dot{s}_k - v_{des})^2, c_a = \sum_k (\dot{s}_k^2 + \dot{d}_k^2), c_{lane} = \sum_k (d_k - d_{des})^2 \quad (34)$$

where v_{des} , y_{des} is the desired longitudinal velocity and lane-offset to be tracked by the ego-vehicle.

The equality constraints $\mathbf{h}(\mathbf{s}^{(q)}, \mathbf{d}^{(q)}) = \mathbf{0}$ from (6) have the following components

$$(s_0, \dot{s}_0, \ddot{s}_0) = (0, v_{x,init}, a_{x,init}) \quad (35a)$$

$$(d_0, \dot{d}_0, \ddot{d}_0) = (d_{init}, v_{y,init}, a_{y,init}, 0) \quad (35b)$$

where the r.h.s in (35a)-(35b) consists of the initial and terminal values of position variables and their derivatives along the axes of the Frenet-Frame.

The inequality constraints $\mathbf{g}(\mathbf{s}^{(q)}, \mathbf{d}^{(q)}) \leq 0$ have the following components.

$$d_{min} \leq d_k \leq d_{max}, \forall k \quad (36a)$$

$$\sqrt{\dot{s}_k^2 + \dot{d}_k^2} \leq v_{max}, \forall k \quad (36b)$$

$$\sqrt{\ddot{s}_k^2 + \ddot{d}_k^2} \leq a_{max}, \forall k \quad (36c)$$

The first inequality (36a) is the lane-boundary constraint. The parameters (d_{min}, d_{max}) are defined with respect to the lane center-line or a reference path. The inequalities (36b)-(36c) bounds the magnitude of the velocity and acceleration respectively.

1) *Benchmarking Environments:* We consider the following collision avoidance scenarios for validating our approach as well as comparing against existing baselines.

a) *Static Obstacles:* We consider a two-lane driving scenario with static obstacles. We consider uncertainty in the location of the obstacles and three different choices for its underlying distribution: Gaussian distribution and Gaussian Mixture Model with two and three modes.

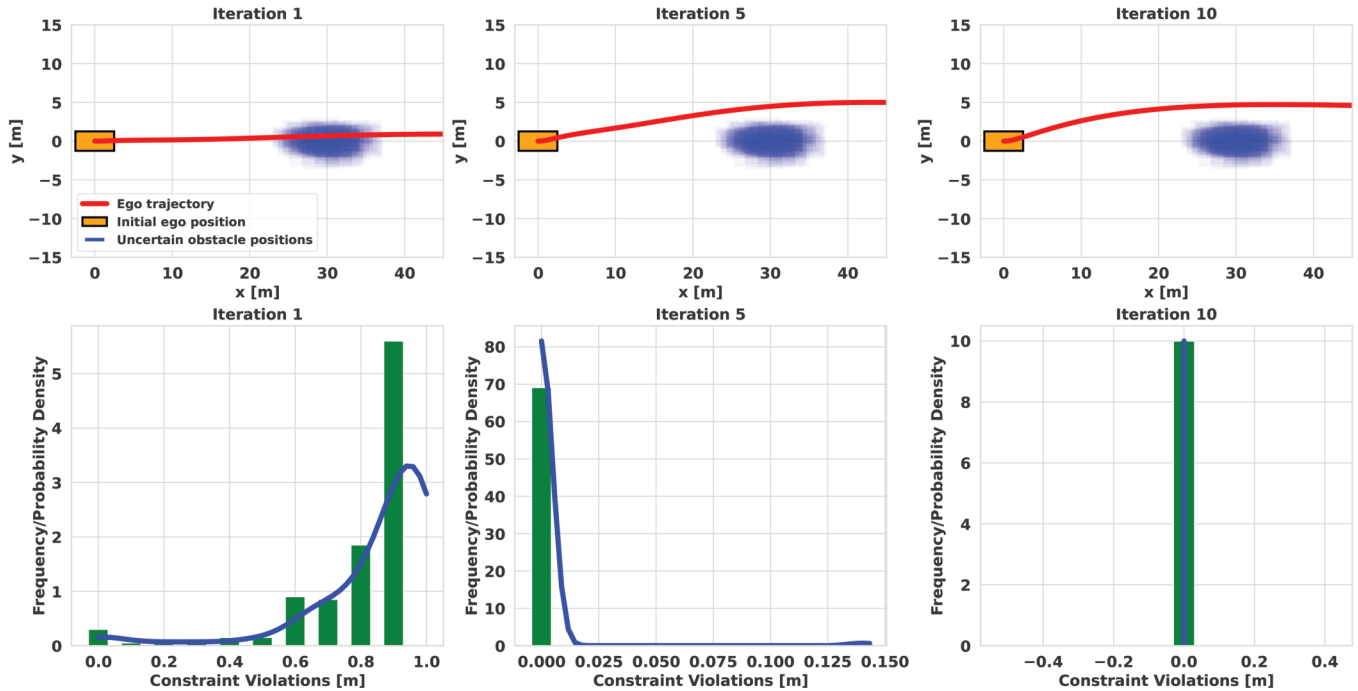


Fig. 8. The top figure shows an ego-vehicle avoiding a uncertain static obstacles. The blue shaded region shows the uncertainty in the location of the obstacle. The trajectory in red depicts the output from Alg. 2 at different iterations. The bottom figure shows the histogram of collision-constraint violation function f across iterations. We also fit a KDE to the histogram to approximately represent $p_{\bar{f}}$. As can be seen, $p_{\bar{f}}$ gradually converges to a Dirac-Delta distribution.

b) Dynamic Environments with Synthetic Trajectory Generator: In this benchmark, we consider a single dynamic obstacle whose future trajectory has two modes representing the different intents in its motions (e.g lane following or change). Moreover, we also consider uncertainty within the execution of each intent. We generate the trajectory distribution by sampling lane and desired velocity set-points from a discrete and a tri-modal GMM respectively, and passing them through the Frenet Space planner described in Appendix A. The scenarios considered cover the typical conflicts observed in autonomous driving.

c) nuScenes Dataset with Trajectron++ Predictor: In this benchmark, we consider the driving scenarios contained in the nuScenes [42] dataset. The scenarios cover a diverse set of conflicts, e.g those arising in un-signalized intersections (see Fig. 11). We use an off-the-shelf pre-trained trajectory predictor Trajectron++ [1] to generate obstacle trajectory distribution. The goal is to generate a safe trajectory for the ego-vehicle that avoids the obstacle trajectory distribution with high probability. Although originally meant for trajectory forecasting, the NuScenes dataset has also been extensively used to validate safe motion planning, including approaches specifically based on risk-aware trajectory optimization [43], [44].

2) Baselines and Benchmarking Protocol: We compare our approach with two baselines obtained by just replacing our r_{MMD}^{emp} with r_{SAA} and r_{CVaR}^{emp} in the trajectory optimizer of Alg. 2. Our aim is to primarily compare the effectiveness of each risk cost at given sample size N' . Thus, to ensure that the planner biases has no influence on the result, we ensure that optimal trajectories resulting from the baselines and our approach

indeed achieve zero risk cost. For example, optimal trajectory resulting from minimizing r_{CVaR}^{emp} would achieve $r_{CVaR}^{emp} = 0$ in all benchmarking examples. By adopting such a protocol, we facilitate explicitly comparing the finite-sample estimation error associated with each risk cost.

We split the sampled obstacle trajectories into two sets namely: optimization and validation set. The former set contains the trajectories from which we construct the r_{MMD}^{emp} (including reduced-set computation), r_{SAA} and r_{CVaR}^{emp} . The validation set trajectories are never seen by the planner. Moreover, the number of trajectories in validation set are at least two-orders of magnitude larger than the optimization set.

For our approach, we sample 100 obstacle trajectories and then select N' out of those as our reduced-set through Alg. 1. We perform collision-checks along only N' to estimate r_{MMD}^{emp} . Since r_{SAA} and r_{CVaR}^{emp} do not have any notion of *reduced-set*, we directly sample N' obstacle trajectories to estimate their values.

3) Metrics: We primarily focus on the collision percentage on the validation set which is computed as follows. We check the computed optimal ego-trajectory for collision with each obstacle trajectory sample in the validation-set and divide the collision-free cases with the size of the validation set. We perform analysis for different sample size (N') of the optimization set. Thus, implicitly N' is also one of our comparison metrics.

B. Qualitative Analysis of Risk-Aware Planning With r_{MMD}^{Emp}

Figure 8 shows the iteration-wise evolution of the optimal trajectory resulting from our optimizer in Alg. 2 while using r_{MMD}^{emp} as the collision risk cost. A static obstacle is placed at an arbitrary position and we introduce uncertainty by




Fig. 9. The figure shows the quantitative analysis (in the form of box plots) of collision rates achieved with different risk costs. We experiment with different types of obstacle trajectory distribution and for each type, we perform analysis across different sample sizes. In the case of Gaussian distribution (a), and $N' = 5$, we observe that r_{MMD}^{emp} leads to a worst-case collision-rate that is two times lower than that achieved using r_{SAA} . The performance gap between r_{MMD}^{emp} and the other two baseline is more enhanced when the underlying obstacle trajectory distribution departs significantly from Gaussian. An example of this is shown in (b) that is based on bimodal obstacle trajectory distribution. We observe that r_{MMD}^{emp} leads to a worst-case collision-rate that is around five times lower than that achieved using r_{CVaR}^{emp}/r_{SAA} .

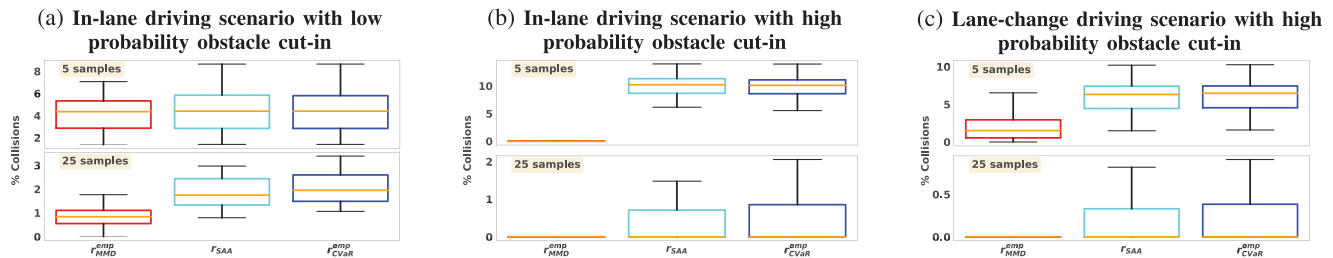


Fig. 10. The box plots show the lower collision-rate achieved by minimizing r_{MMD}^{emp} over the baseline risk cost r_{SAA} and r_{CVaR}^{emp} for the dynamic obstacle avoidance scenario involving multi-modal obstacle trajectories.

adding noise to its ground-truth position. As the iterations progress, the optimal trajectory shifts away from the potential obstacle locations making the trajectory increasingly safe. In the histogram plot of constraint violations, we observe that with each iteration, the probability of constraint violations goes down and by the 10th iteration, the distribution $p_{\bar{f}}$ of the collision constraint violation function \bar{f} has already converged to the Dirac delta distribution p_{δ} .

C. Benchmarking in Synthetic Static Environments

In this benchmark, we randomly generated, 100 different obstacle configurations, each with 3 static obstacles. We introduced noise in the positions of all the obstacles. Three different types of noises were used: Gaussian, Bimodal GMM and Trimodal GMM. The initial position of the ego vehicle was fixed at the same value for all experiments. Similarly, the initial velocity of the ego vehicle was chosen as 3m/s in every experiment. We report the collision statistics (Number of collisions of the optimal trajectory with the obstacle samples from the validation set (recall Sections VI-A2 and VI-A3)).

We present the qualitative results in the complimentary video. The statistical benchmarking is presented in Fig. 9. The following points are worth noting. First, the r_{MMD}^{emp} trajectories show far lower number of collisions on the validation set than that computed using r_{CVaR}^{emp} and r_{SAA} risk costs at all sample sizes. For example, under Gaussian noise, r_{MMD}^{emp} leads to a worst-case collision-rate that is two times lower than that achieved using r_{SAA} . Second, the difference in performance between r_{MMD}^{emp} and r_{CVaR}^{emp}/r_{SAA} becomes more stark as the underlying obstacle trajectories become increasingly more multi-modal. This showcases the ability of r_{MMD}^{emp} to leverage RKHS embedding of distributions to capture the underlying

probability distribution of obstacle trajectories and collision constraint residual function $p_{\bar{f}}$ with very few samples. Finally, r_{MMD}^{emp} leads to fastest reduction in collision-rate with increase in sample size.

D. Benchmarking in Synthetic Dynamic Environments

In this benchmark, we introduced a single dynamic obstacle with uncertain multi-modal future trajectory. The quantitative results are summarized in Fig. 10a-10c. We once again observe that optimal trajectories obtained by minimizing r_{MMD}^{emp} provides the best collision-avoidance performance across all sample-size N' . In particular, while using $N' = 5$ obstacle trajectory samples, the median collision-rate achieved by minimizing r_{MMD}^{emp} is 3 times lower than that resulting from r_{CVaR}^{emp} and r_{SAA} . The performance difference shown in Fig. 10b, is even more stark with r_{MMD}^{emp} demonstrating zero collisions on the validation set.

E. Benchmarking on nuScenes Dataset Using Trajectron++ Predictor

We evaluated a total of 160 scenes in this dataset. Each scene comes with a reference center-line and a designated ego vehicle. Any vehicle other than the ego was treated as an obstacle. We considered two settings namely, in-lane driving and lane-change. In the former, the ego-vehicle needs to avoid collisions by just changing the speed, while the latter allows for the full set of maneuvers to the ego-vehicle. We report the collision statistics in Fig. 12. We used Trajectron++ [1] as a deep neural network based trajectory predictor. We observed that for both scenarios (in-lane and lane-change), our approach based on r_{MMD}^{emp} has the lowest worst-case collision-rate among




Fig. 11. Risk-Aware trajectory planning in an unprotected intersection. The blue trajectories show the trajectory predictions of the neighboring vehicles. As can be seen, the predictions are highly multi-modal capturing different driving intents. Typical ego-vehicle trajectories resulting from different collision risk costs are also shown. It can be seen that r_{MMD}^{emp} correctly understands that the cyan-vehicle is more likely to turn left. Hence, it shows least deviation from the recorded ground-truth (black), human-driven trajectory in the dataset. In contrast, r_{CVaR}^{emp} and r_{SAA} incorrectly puts more emphasis on the less likely scenario of cyan vehicle turning right. As a result, the resulting trajectories show a more deviation from the ground-truth trajectory.

all the approaches. The improvement is approximately two-times in the lane-change scenario when using $N' = 25$ samples from obstacle trajectory distribution. The median collision-rate for our approach and the baselines were zero.

A typical example of trajectories obtained in this benchmark is shown in Fig. 11 which demonstrates an unprotected intersection. In this setting, the ego-vehicle must safely navigate through cross-traffic while avoiding dynamic obstacles. Unlike standard intersection maneuvers with dedicated traffic signals or stop signs, an unprotected intersection requires the ego-vehicle to make real-time decisions based on surrounding vehicle behavior. The primary challenge in this setting is accurately estimating the risk of collision while ensuring smooth and efficient trajectory execution. As can be seen in Fig. 11 the ego-vehicle approaches the intersection while multiple obstacles follow stochastic trajectories across its path. The predicted trajectories cover multi-modal intents (e.g turn left or right). The trajectories resulting from minimizing r_{MMD}^{emp} is closest to the ground-truth (human-driven) recorded trajectory in the same scenario. This in turn can be attributed to the ability to r_{MMD}^{emp} to correctly identify (through its *reduced-set*, recall Fig. 7) the more likely maneuvers. For example, consider the cyan obstacle/vehicle. The predicted trajectories cover both left and right turns. But the ground-truth motion for the obstacle is the left turn maneuver and r_{MMD}^{emp} is able to correctly give more importance to left-turn predictions. In contrast, the r_{SAA} , r_{CVaR}^{emp} show overly conservative maneuver driven by the trajectory samples that predict right-turn for the cyan obstacle.

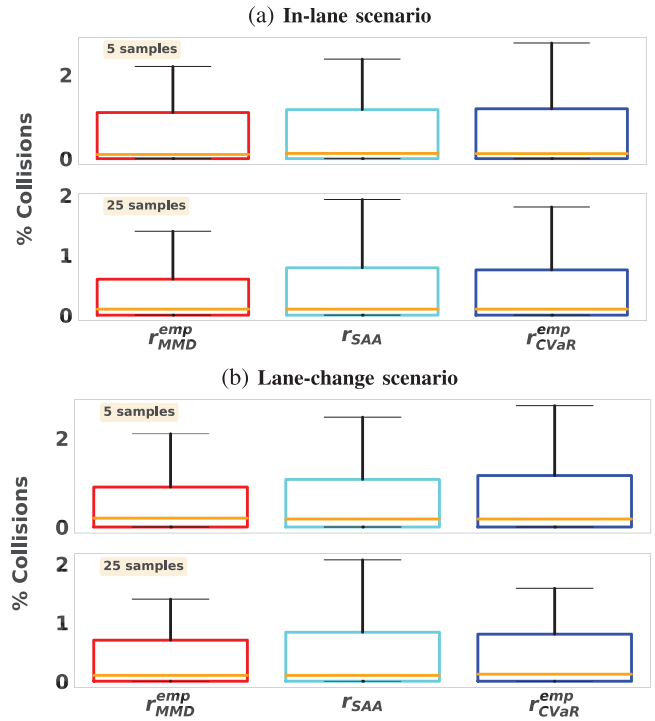


Fig. 12. Comparing how well different risks r_{MMD}^{emp} , r_{SAA} , r_{CVaR}^{emp} perform on real-world datasets with Neural Network based trajectory predictors (Trajectron++ [1] in this case). Typically, these predictors define a very complex distribution over the possible trajectories of the neighboring vehicles (obstacles). Hence, it is challenging to estimate the collision risk with a few predicted samples drawn from the distribution. As can be seen, our surrogate r_{MMD}^{emp} leads to optimal trajectories with lowest collision-rate in both in-lane driving and lane-change scenarios.

F. Ablation Results

Fig. 13 performs an ablation when instead of using Alg. 1, we form the *reduced-set* by random sub-selection (recall Section IV-C1). It can be seen that constructing r_{MMD}^{emp} on the optimal *reduced-set* leads to lower collision-rate across all scenes. The difference between the random and the optimal variant is particularly stark when the underlying trajectory distribution is much more complex than a Gaussian. Moreover, the variance in the collision-rate across the benchmarking scenes is lower when using the optimal *reduced-set*.

G. Comparison With a Scenario Approach

Fig. 14 presents an ablation study in a dynamic obstacle avoidance setting, comparing our approach to the classical scenario approach for risk-aware planning. This approach boils down to solving the following deterministic problem

$$\min_{\mathbf{s}, \mathbf{d}} c(\mathbf{s}^{(q)}, \mathbf{d}^{(q)}) \quad (37)$$

$$\mathbf{h}(\mathbf{s}^{(q)}, \mathbf{d}^{(q)}) = \mathbf{0} \quad (38)$$

$$\mathbf{g}(\mathbf{s}^{(q)}, \mathbf{d}^{(q)}) \leq \mathbf{0} \quad (39)$$

$$f_k(s_k, d_k, \mathbf{j}\boldsymbol{\tau}_k) \leq 0, \forall k, j \quad (40)$$

The inequality constraints (40) enforces a deterministic collision avoidance condition for the j^{th} sampled obstacle motions scenario $\mathbf{j}\boldsymbol{\tau}_k$ (recall Section II-C). We solve the scenario approach with CasADi [45]. For a fair comparison,

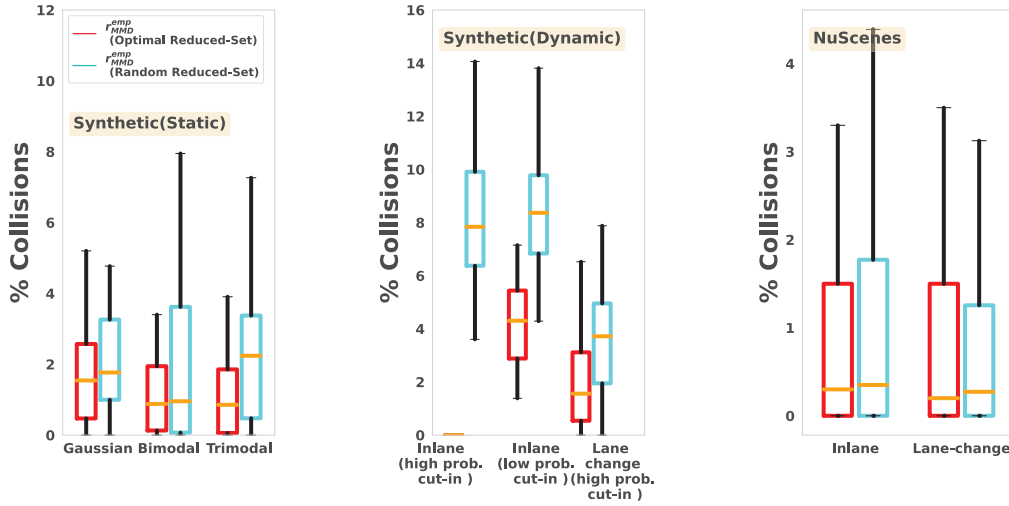


Fig. 13. Comparison of collision-rate achieved by minimizing r_{MMD}^{emp} constructed with the optimal and random *reduced-set* of size $N' = 5$. The former set-up leads to safer trajectories with a significantly better median and worst-case statistics. The difference is particularly stark when the obstacle trajectory distribution departs significantly from a Gaussian.

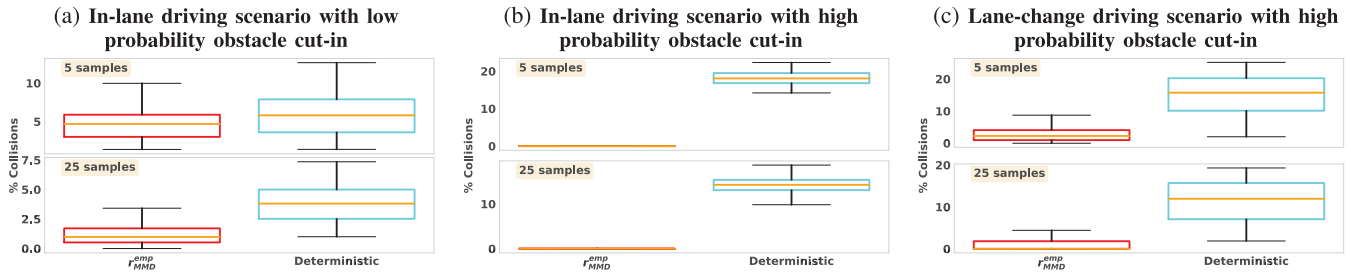


Fig. 14. The box plots show the lower collision-rate achieved by minimizing r_{MMD}^{emp} over the noise-ignorant deterministic baseline for the synthetic dynamic obstacle avoidance setting, involving multi-modal obstacle trajectories.

we impose as many obstacle avoidance constraints as the *reduced-set* samples used in minimizing r_{MMD}^{emp} . The exact mathematical form of the cost and constraints is provided in Section VI. Fig. 14 demonstrates that minimizing r_{MMD}^{emp} achieves significantly lower collision rates than the scenario approach, underscoring the effectiveness of our approach in handling multi-modal uncertainty, improving collision avoidance, and enhancing planning robustness.

H. Impact of Using an Optimal Reduced-Set for Baselines

The baselines r_{CVaR}^{emp} and r_{SAA} do not incorporate the concept of a *reduced-set*, which is mathematically well-founded only for r_{MMD}^{emp} . This distinction arises due to the kernel-based distributional properties of MMD, as discussed in [20]. To the best of our knowledge, no equivalent *reduced-set* formulation exists for r_{CVaR}^{emp} or r_{SAA} . However, for the purpose of fairness in comparison, we conducted an additional ablation study on the nuScenes dataset. In this experiment, we applied the optimal *reduced-set* of obstacle trajectories—originally derived for r_{MMD}^{emp} —to the computation of r_{CVaR}^{emp} and r_{SAA} . The results, presented in Figure 15, confirm that the introduction of the *reduced-set* does not provide a significant advantage to these baselines. This outcome is expected, as r_{CVaR}^{emp} and r_{SAA} rely on independent random sampling from the underlying obstacle trajectory distribution, rather than optimizing over a subset

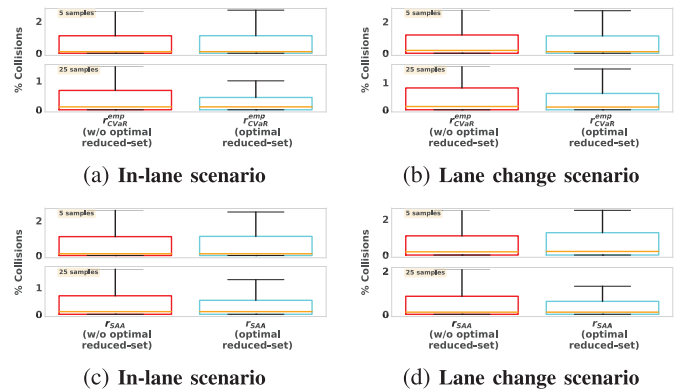


Fig. 15. Ablation study on the nuScenes dataset: Comparison of performance of r_{CVaR}^{emp} and r_{SAA} with and without the use of optimal *reduced-set* of obstacle trajectories. The results show that the collision-avoidance performance of the baselines r_{CVaR}^{emp} and r_{SAA} remains largely unaffected with the introduction of the optimal *reduced-set*, reinforcing the theoretical distinction between MMD-based risk estimation and conventional risk metrics.

of the samples in a manner mathematically consistent with r_{MMD}^{emp} . Consequently, while r_{MMD}^{emp} is bound to benefit from the *reduced-set* formulation, the baselines remain largely unaffected by this modification. This experiment further highlights the theoretical distinction between MMD-based risk estimation and conventional risk metrics such as CVaR and SAA.

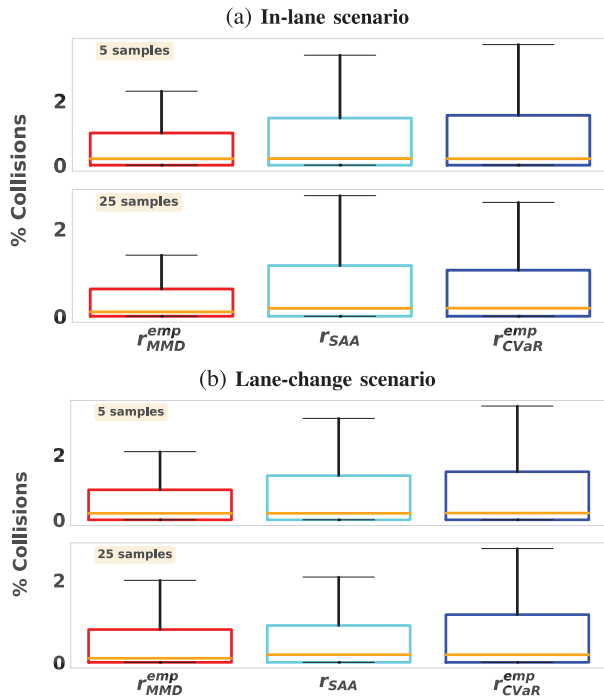


Fig. 16. Dynamics bicycle model.

TABLE I
COMPUTATION TIME (IN SECONDS) VS NUMBER OF
REDUCED-SET SAMPLES

Number of Reduced-set samples	Number of obstacles		
	1 obs. (Alg. 1/Alg.2)	2 obs. (Alg. 1/Alg.2)	3 obs. (Alg. 1/Alg.2)
5 samples	0.005/0.01	0.01/0.01	0.016/0.01
10 samples	0.006/0.01	0.012/0.01	0.019/0.01
20 samples	0.006/0.01	0.012/0.01	0.017/0.01
30 samples	0.007/0.01	0.014/0.01	0.02/0.01
40 samples	0.007/0.01	0.014/0.01	0.02/0.01
50 samples	0.008/0.01	0.015/0.01	0.023/0.01

I. Risk-Aware Planning With Dynamics Model

Our proposed risk surrogate r_{MMD}^{emp} just requires the information about the ego-vehicle and the obstacle motions. Thus, it is agnostic to the actual motion model of the ego-vehicle. The results shown in the previous subsection assumed kinematic bi-cycle model. However, we can get similar results even if we use a more complex dynamics for the ego-vehicle that includes tire forces and side-slips. To further validate this assertion, we adapted our Alg. 2 to work with the dynamics model derived from [46], [47]. For comparison, we also minimized r_{CVaR}^{emp} and r_{SAA} under the same dynamics.

The results are summarized in Fig. 16 for the nuScenes dataset with Trajectron++ as the trajectory predictor. As can be seen, r_{MMD}^{emp} still outperforms r_{CVaR}^{emp} and r_{SAA} , even when the underlying motion model is more complex. The pattern is very similar to what we observed in Fig. 12 for the bi-cycle kinematic model.

J. Computation Time

Table I shows the computation time required for our reduced-set optimization Alg. 1 and Alg. 2 for different sample

sizes on an RTX 3090 Desktop with 16 GB RAM. We report the timing separately for Alg. 1 and 2. The former depends on the number of obstacles and also on the size of the optimal *reduced-set* (N') that we want to construct. At any given N' , Alg. 1 show almost linear growth in the computation-time with respect to the number of obstacles. However, once Alg. 1 is done, the subsequent computation of estimating r_{MMD}^{emp} for a given ego-vehicle trajectory can be suitably parallelized. Hence the increase in computation-time for Alg. 2 with respect to the number of obstacles is very minimal. The run-time indicates that our approach can be used for fast re-planning in a receding horizon fashion.

VII. CONCLUSION AND FUTURE WORK

Estimating and minimizing collision risk based on predicted obstacle trajectories is one of the fundamental problems in autonomous driving. Existing works have used techniques like SAA (r_{SAA} , Eqn. 11) and CVaR (r_{CVaR}^{emp} , Eqn. 12) to estimate the collision constraint residual function (recall (10)) and define a risk surrogate conditioned on the ego-vehicle and obstacle trajectories. Although effective, a key limitation of these techniques is the lack of a systematic approach to maximize their effectiveness at a given finite sample size. To address this gap, we proposed using the MMD between the distribution of collision-constraint residuals and the Dirac-Delta distribution to model collision risk (r_{MMD}^{emp} , Eqn. 23).

The choice of our risk surrogate allows us to use the notion of *reduced-set* to maximize its effectiveness at a given sample size. In other words, the *reduced-set* enables selecting only a subset of predicted obstacle trajectories to formulate our risk surrogate while minimizing accuracy loss. However, fully leveraging its potential requires moving beyond the conventional approach of constructing *reduced-sets* through random sub-selection [11]. Motivated by this, we proposed a novel bi-level optimization method to construct *reduced-sets* while ensuring minimal loss in downstream risk estimation. As a byproduct, the bi-level optimizer also provides an optimal kernel parameter.

We compared the effectiveness of r_{MMD}^{emp} , r_{SAA} , and r_{CVaR}^{emp} when used as a risk-cost in trajectory optimization. The results show that r_{MMD}^{emp} significantly improves performance, particularly at lower sample sizes. This improvement is especially evident when the underlying prediction distribution deviates substantially from a Gaussian. Our findings demonstrate the suitability of r_{MMD}^{emp} for use alongside large neural network-based predictors, which induce unknown or intractable distributions over obstacle trajectories. We also provided preliminary results on driving scenes from the nuScenes dataset [42] using Trajectron++ as our trajectory predictor.

MMD-OPT is designed for real-time applicability while improving sample efficiency for risk-aware planning. The *reduced-set* selection mechanism keeps computational costs manageable, and our GPU-accelerated implementation ensures practical feasibility. Validations on nuScenes with state-of-the-art trajectory prediction models further demonstrate MMD-OPT's effectiveness in handling dynamic, multi-obstacle interactions. Though our current formulation just covers an open-loop optimal planning, it can be trivially extended to

a Model Predictive Control (MPC) setting by invoking the planner (Alg. 2) at each control cycle. The fast computation-times shown in Table I indicate that the resulting MPC can run at around 50 Hz.

Our approach is not limited to autonomous driving, and future efforts will focus on extending it to 3D drone navigation and manipulation in cluttered environments. We also aim to extend our MMD-based risk surrogate to model collision risk under uncertain vehicle dynamics. Future work will also focus on optimizing MMD-OPT for embedded systems and extending its applicability to dense urban traffic environments.

APPENDIX

A. Frenet Planner

Let $\mathbf{b} = (b_d, b_v)$ be the lateral offset and desired velocity setpoints. Our Frenet planner inspired from [25] boils down to solving the following trajectory optimization.

$$\min \sum_k c_s + c_l + c_v \quad (41a)$$

$$\mathbf{h}(\mathbf{s}^{(q)}, \mathbf{d}^{(q)}) = \mathbf{0} \quad (41b)$$

$$c_s(\dot{s}_k, \ddot{d}_k) = \dot{s}_k^2 + \ddot{d}_k^2 \quad (42a)$$

$$c_l(d_k, \dot{d}_k, \ddot{d}_k) = (\ddot{d}_k - \kappa_p(d_k - b_d) - \kappa_v \dot{d}_k)^2 \quad (42b)$$

$$c_v(\dot{s}_k, \ddot{s}_k) = (\ddot{s}_k - \kappa_p(\dot{s}_k - b_v))^2 \quad (42c)$$

The first term $c_s(\cdot)$ in the cost function (41a) ensures smoothness in the planned trajectory by penalizing high accelerations at discrete time instants. The last two terms ($c_l(\cdot), c_v(\cdot)$) model the tracking of lateral offset (b_d) and forward velocity (b_v) set-points respectively with gain (κ_p, κ_v). The above optimization is an equality constrained QP that have a closed-form solution.

B. Simplification of r_{Mmd}^{Emp} Based On Kernel Trick

We expand equation (23) in the following manner

$$\begin{aligned} \|\widehat{\mu}[\overline{f'}] - \widehat{\mu}[\delta]\|_{\mathcal{H}}^2 &= \langle \widehat{\mu}[\overline{f'}], \widehat{\mu}[\overline{f'}] \rangle - 2\langle \widehat{\mu}[\overline{f'}], \widehat{\mu}[\delta] \rangle \\ &\quad + \langle \widehat{\mu}[\delta], \widehat{\mu}[\delta] \rangle \end{aligned} \quad (43)$$

Substituting RKHS embeddings of $\overline{f'}$ (25) and $\delta \sim p_\delta$ in (43)

$$\langle \widehat{\mu}[\overline{f'}], \widehat{\mu}[\overline{f'}] \rangle = \left\langle \sum_{i=1}^{N'} i \beta K(i\overline{f'}, \cdot), \sum_{j=1}^{N'} j \beta K(j\overline{f'}, \cdot) \right\rangle \quad (44a)$$

$$\langle \widehat{\mu}[\overline{f'}], \widehat{\mu}[\delta] \rangle = \left\langle \sum_{i=1}^{N'} i \beta K(i\overline{f'}, \cdot), \frac{1}{N'} \sum_{j=1}^{N'} K(j\delta, \cdot) \right\rangle \quad (44b)$$

$$\langle \widehat{\mu}[\delta], \widehat{\mu}[\delta] \rangle = \left\langle \frac{1}{N'} \sum_{i=1}^{N'} K(i\delta, \cdot), \frac{1}{N'} \sum_{j=1}^{N'} K(j\delta, \cdot) \right\rangle \quad (44c)$$

where the computations are done over the reduced-set introduced in Section IV-C. Substituting (44) in (43) and using the kernel trick introduced in Section III-A, we can now derive an optimization friendly expression for the r.h.s of (23) as follows

$$\|\widehat{\mu}[\overline{f'}] - \widehat{\mu}[\delta]\|_2^2 = \mathbf{C}_\beta \mathbf{K}_{\overline{f'}, \overline{f'}} \mathbf{C}_\beta^T - 2\mathbf{C}_\beta \mathbf{K}_{\overline{f'}, \delta} \mathbf{C}_\beta^T + \mathbf{C}_\beta \mathbf{K}_{\delta, \delta} \mathbf{C}_\beta^T \quad (45)$$

where \mathbf{C}_β and $\mathbf{C}_{\overline{\beta}}$ are weight vectors given by

$$\mathbf{C}_\beta = [{}^1\beta, {}^2\beta, \dots, {}^{N'}\beta], \mathbf{C}_{\overline{\beta}} = \left[\frac{1}{N'}, \frac{1}{N'}, \dots, \frac{1}{N'} \right]$$

and $\mathbf{K}_{\overline{f'}, \overline{f'}}$, $\mathbf{K}_{\overline{f'}, \delta}$ and $\mathbf{K}_{\delta, \delta}$ are kernel matrices defined as

$$\mathbf{K}_{\overline{f'}, \overline{f'}} = \begin{bmatrix} K({}^1\overline{f'}, {}^1\overline{f'}) & \dots & K({}^1\overline{f'}, {}^{N'}\overline{f'}) \\ K({}^2\overline{f'}, {}^1\overline{f'}) & \dots & K({}^2\overline{f'}, {}^{N'}\overline{f'}) \\ \vdots & \ddots & \vdots \\ K({}^{N'}\overline{f'}, {}^1\overline{f'}) & \dots & K({}^{N'}\overline{f'}, {}^{N'}\overline{f'}) \end{bmatrix}$$

$$\mathbf{K}_{\overline{f'}, \delta} = \begin{bmatrix} K({}^1\overline{f'}, 0) & \dots & K({}^1\overline{f'}, 0) \\ K({}^2\overline{f'}, 0) & \dots & K({}^2\overline{f'}, 0) \\ \vdots & \ddots & \vdots \\ K({}^{N'}\overline{f'}, 0) & \dots & K({}^{N'}\overline{f'}, 0) \end{bmatrix}, \mathbf{K}_{\delta, \delta} = \mathbf{1}_{N' \times N'}$$

REFERENCES

- [1] B. Ivanovic and M. Pavone, "The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 2375–2384.
- [2] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social GAN: Socially acceptable trajectories with generative adversarial networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 2255–2264.
- [3] N. Lee, W. Choi, P. Vernaza, C. B. Choy, P. H. S. Torr, and M. Chandraker, "DESIRE: Distant future prediction in dynamic scenes with interacting agents," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 336–345.
- [4] T. Mitchell Roddenberry, N. Glaze, and S. Segarra, "Principled simplicial neural networks for trajectory prediction," 2021, *arXiv:2102.10058*.
- [5] S. Zamboni, Z. T. Kefato, S. Girdzijauskas, C. Norén, and L. Dal Col, "Pedestrian trajectory prediction with convolutional neural networks," *Pattern Recognit.*, vol. 121, Jan. 2022, Art. no. 108252. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0031320321004325>
- [6] C. Xu, M. Li, Z. Ni, Y. Zhang, and S. Chen, "GroupNet: Multiscale hypergraph neural networks for trajectory prediction with relational reasoning," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 6488–6497.
- [7] H. Zhu and J. Alonso-Mora, "Chance-constrained collision avoidance for MAVs in dynamic environments," *IEEE Robot. Autom. Lett.*, vol. 4, no. 2, pp. 776–783, Apr. 2019.
- [8] L. Blackmore, M. Ono, and B. C. Williams, "Chance-constrained optimal path planning with obstacles," *IEEE Trans. Robot.*, vol. 27, no. 6, pp. 1080–1094, Dec. 2011.
- [9] M. Castillo-Lopez, P. Ludvig, S. A. Sajadi-Alamdari, J. L. Sanchez-Lopez, M. A. Olivares-Mendez, and H. Voos, "A real-time approach for chance-constrained motion planning with dynamic obstacles," *IEEE Robot. Autom. Lett.*, vol. 5, no. 2, pp. 3620–3625, Apr. 2020.
- [10] C. Dawson, A. Jasour, A. Hofmann, and B. Williams, "Provably safe trajectory optimization in the presence of uncertain convex obstacles," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2020, pp. 6237–6244.
- [11] C.-J. Simon-Gabriel, A. Scibior, I. Tolstikhin, and B. Schölkopf, "Consistent kernel mean estimation for functions of random variables," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 29, Nov. 2016, pp. 1732–1740.
- [12] A. Gretton, K. Borgwardt, M. J. Rasch, B. Schölkopf, and A. J. Smola, "A kernel two-sample test," *J. Mach. Learn. Res.*, vol. 13, no. 1, pp. 723–773, Mar. 2012.
- [13] B. K. Pagnoncelli, S. Ahmed, and A. Shapiro, "Sample average approximation method for chance constrained programming: Theory and applications," *J. Optim. Theory Appl.*, vol. 142, no. 2, pp. 399–416, Mar. 2009.
- [14] J. Yin, Z. Zhang, and P. Tsiotras, "Risk-aware model predictive path integral control using conditional value-at-risk," 2022, *arXiv:2209.12842*.
- [15] A. Liniger and L. van Gool, "Safe motion planning for autonomous driving using an adversarial road model," 2020, *arXiv:2005.07691*.
- [16] 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," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2022, pp. 11087–11093.

- [17] Y. Shirai, D. K. Jha, and A. U. Raghunathan, "Covariance steering for uncertain contact-rich systems," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2023, pp. 7923–7929, doi: [10.1109/ICRA48891.2023.10160249](https://doi.org/10.1109/ICRA48891.2023.10160249).
- [18] T. Lew, R. Bonalli, and M. Pavone, "Risk-averse trajectory optimization via sample average approximation," *IEEE Robot. Autom. Lett.*, vol. 9, no. 2, pp. 1500–1507, Feb. 2024, doi: [10.1109/LRA.2023.3331889](https://doi.org/10.1109/LRA.2023.3331889).
- [19] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming: Modeling and Theory* (MOS-SIAM Series on Optimization), 2nd ed., Philadelphia, PA, USA: SIAM, 2014. [Online]. Available: <https://books.google.ee/books?id=ob4ABAAAQBAJ>
- [20] K. Muandet, K. Fukumizu, B. K. Sriperumbudur, and B. Schölkopf, "Kernel mean embedding of distributions: A review and beyond," *Found. Trends Mach. Learn.*, vol. 10, nos. 1–2, pp. 1–141, Jan. 2017, doi: [10.1561/9781680832891](https://doi.org/10.1561/9781680832891).
- [21] A. Gretton, K. Borgwardt, M. J. Rasch, B. Schölkopf, and A. J. Smola, "A kernel method for the two-sample problem," in *Proc. Adv. Neural Inf. Process. Syst.*, Sep. 2007, pp. 513–520. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2006/file/e9fb2eda3d9c55a0d89c98d6c54b5b3e-Paper.pdf
- [22] Z. I. Botev, D. P. Kroese, R. Y. Rubinstein, and P. L'Ecuyer, "The cross-entropy method for optimization," in *Handbook of Statistics*, vol. 31. Amsterdam, The Netherlands: Elsevier, 2013, pp. 35–59.
- [23] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Aggressive driving with model predictive path integral control," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, Stockholm, Sweden, May 2016, pp. 1433–1440.
- [24] M. Bhardwaj et al., "Storm: An integrated framework for fast joint-space model-predictive control for reactive manipulation," in *Proc. 5th Conf. Robot Learn. (CoRL)*, vol. 164, 2022, pp. 750–759.
- [25] J. Wei, J. M. Snider, T. Gu, J. M. Dolan, and B. Litkouhi, "A behavioral planning framework for autonomous driving," in *Proc. IEEE Intell. Vehicles Symp.*, Jun. 2014, pp. 458–464.
- [26] H. Masnavi, J. Shrestha, M. Mishra, P. B. Sujit, K. Kruusamäe, and A. K. Singh, "Visibility-aware navigation with batch projection augmented cross-entropy method over a learned occlusion cost," *IEEE Robot. Autom. Lett.*, vol. 7, no. 4, pp. 9366–9373, Oct. 2022.
- [27] V. K. Adajania, A. Sharma, A. Gupta, H. Masnavi, K. M. Krishna, and A. K. Singh, "Multi-modal model predictive control through batch non-holonomic trajectory optimization: Application to highway driving," *IEEE Robot. Autom. Lett.*, vol. 7, no. 2, pp. 4220–4227, Apr. 2022.
- [28] Y. Rasekhipour, A. Khajepour, S.-K. Chen, and B. Litkouhi, "A potential field-based model predictive path-planning controller for autonomous road vehicles," *IEEE Trans. Intell. Transp. Syst.*, vol. 18, no. 5, pp. 1255–1267, May 2017.
- [29] W. Ding, L. Zhang, J. Chen, and S. Shen, "Safe trajectory generation for complex urban environments using spatio-temporal semantic corridor," *IEEE Robot. Autom. Lett.*, vol. 4, no. 3, pp. 2997–3004, Jul. 2019.
- [30] G. C. Calafiore and M. C. Campi, "The scenario approach to robust control design," *IEEE Trans. Autom. Control*, vol. 51, no. 5, pp. 742–753, May 2006.
- [31] G. C. Calafiore and L. Fagiano, "Robust model predictive control via scenario optimization," *IEEE Trans. Autom. Control*, vol. 58, no. 1, pp. 219–224, Jan. 2013.
- [32] H. Ahn, C. Chen, I. M. Mitchell, and M. Kamgarpour, "Safe motion planning against multimodal distributions based on a scenario approach," *IEEE Control Syst. Lett.*, vol. 6, pp. 1142–1147, 2022.
- [33] D. Bernardini and A. Bemporad, "Scenario-based model predictive control of stochastic constrained linear systems," in *Proc. 48th IEEE Conf. Decis. Control (CDC)*, Dec. 2009, pp. 6333–6338.
- [34] G. Schildbach, L. Fagiano, C. Frei, and M. Morari, "The scenario approach for stochastic model predictive control with bounds on closed-loop constraint violations," *Automatica*, vol. 50, no. 12, pp. 3009–3018, Dec. 2014. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0005109814004166>
- [35] J.-J. Zhu, M. Diehl, and B. Schölkopf, "A kernel mean embedding approach to reducing conservativeness in stochastic programming and control," in *Proc. Learn. Dyn. Control*, Jan. 2020, pp. 915–923.
- [36] K. A. Mustafa, D. J. Ornia, J. Kober, and J. Alonso-Mora, "RACP: Risk-aware contingency planning with multi-modal predictions," *IEEE Trans. Intell. Vehicles*, early access, Jun. 7, 2024, doi: [10.1109/TIV.2024.3411530](https://doi.org/10.1109/TIV.2024.3411530).
- [37] H. D. Nguyen, M. Choi, and K. Han, "Risk-informed decision-making and control strategies for autonomous vehicles in emergency situations," *Accident Anal. Prevention*, vol. 193, Sep. 2023, Art. no. 107305.
- [38] Y. Nemmour, H. Kremer, B. Schölkopf, and J.-J. Zhu, "Maximum mean discrepancy distributionally robust nonlinear chance-constrained optimization with finite-sample guarantee," in *Proc. IEEE 61st Conf. Decis. Control*, Jul. 2022, pp. 5660–5667.
- [39] B. Sharma, A. Sharma, K. M. Krishna, and A. K. Singh, "Hilbert space embedding-based trajectory optimization for multi-modal uncertain obstacle trajectory prediction," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2023, pp. 7448–7455.
- [40] B. Gopalakrishnan, A. K. Singh, K. M. Krishna, and D. Manocha, "Solving chance-constrained optimization under nonparametric uncertainty through Hilbert space embedding," *IEEE Trans. Control Syst. Technol.*, vol. 30, no. 3, pp. 901–916, May 2022.
- [41] J. Bradbury et al. (2018). *JAX: Composable Transformations of Python+NumPy Programs*. [Online]. Available: <http://github.com/google/jax>
- [42] H. Caesar et al., "nuScenes: A multimodal dataset for autonomous driving," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 11621–11631.
- [43] P. Karkus, B. Ivanovic, S. Mannor, and M. Pavone, "DiffStack: A differentiable and modular control stack for autonomous vehicles," in *Proc. Conf. Robot Learn.*, Jan. 2022, pp. 2170–2180.
- [44] K. Ren, H. Ahn, and M. Kamgarpour, "Chance-constrained trajectory planning with multimodal environmental uncertainty," *IEEE Control Syst. Lett.*, vol. 7, pp. 13–18, 2023.
- [45] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi: A software framework for nonlinear optimization and optimal control," *Math. Program. Comput.*, vol. 11, no. 1, pp. 1–36, Mar. 2019.
- [46] J. Becker, N. Imholz, L. Schwarzenbach, E. Ghignone, N. Baumann, and M. Magno, "Model(-) and acceleration-based pursuit controller for high-performance autonomous racing," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2023, pp. 5276–5283, doi: [10.1109/ICRA48891.2023.10161472](https://doi.org/10.1109/ICRA48891.2023.10161472).
- [47] J. L. Vázquez, M. Bruhlmeier, A. Liniger, A. Rupenyan, and J. Lygeros, "Optimization-based hierarchical motion planning for autonomous racing," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Oct. 2020, pp. 2397–2403.



Basant Sharma (Graduate Student Member, IEEE) received the Bachelor of Engineering (Hons.) degree in mechanical engineering from Birla Institute of Technology and Science, Pilani, K.K. Birla Goa Campus, India, and the Master of Science degree in mathematics from IIT Madras, India, in 2013 and 2020, respectively. He is currently pursuing the Ph.D. degree in computer engineering with the University of Tartu, Estonia. His primary research interests include stochastic trajectory optimization (with a focus on chance-constrained optimization), and motion planning and control.



Arun Kumar Singh received the Ph.D. degree from IIIT-Hyderabad, India, in December 2014. He is currently an Associate Professor at the University of Tartu, Estonia. His research interests lie in developing optimization algorithms customized for robotic applications, embedding classical algorithmic and control theoretic reasoning into neural network pipelines, and planning under uncertainty. His works span a diverse application domain ranging from manipulation, drone navigation to autonomous driving.