Sample Efficient Risk Estimation Through Maximum Mean Discrepancy

Basant Sharma, Arun Kumar Singh

***Abstract*—For an autonomous vehicle to operate safely, it must be able to accurately estimate the collision risk of any planned trajectory, accounting for uncertainty in its own motion and in the environment. The challenge is that when this uncertainty is complex and non-Gaussian, we can no longer rely on a simple analytical formula to calculate risk. This forces us to use sampling-based methods, where sample efficiency becomes paramount. We introduce our novel approach to this problem, which leverages Distribution Embedding in a Reproducing Kernel Hilbert Space and the Maximum Mean Discrepancy (MMD) to create a highly sample-efficient risk estimator. We have validated this technique in two critical applications: first, planning around dynamic obstacles with multi-modal future predictions, and second, stochastic Model Predictive Control under uncertain dynamics. Our results consistently show that this MMD-based method is significantly more sample-efficient than the widely-used Conditional Value at Risk (CVaR) metric.**

I. Introduction

Collision avoidance is central to autonomous driving, requiring the ego vehicle to plan safe trajectories under uncertainty in both environment and dynamics. Such uncertainty, often complex and non-Gaussian, cannot be captured by simple parametric models, making sampling-based risk estimation necessary. In this setting, computational cost and sample efficiency are critical.

We study two applications. First, with deterministic ego dynamics, uncertainty arises from neighboring vehicles whose predicted trajectories are multi-modal [1]–[3]. Our approach [4] exploits these predictions to distinguish high- from low-probability behaviors, enabling safer planning. Second, with deterministic obstacle motion, uncertainty stems from noise in ego vehicle's dynamics and localization errors. Here, risk-aware trajectory optimization provides a principled framework [5], [6], but analytic models are intractable and Gaussian approximations unreliable [7], [8].

Contribution: Our work [4], [9] proposes a sample-efficient surrogate for collision risk estimation that is robust to arbitrary uncertainties in both environmental and vehicle dynamics. The method is founded on the concept of distribution embedding in a Reproducing Kernel Hilbert Space (RKHS), utilizing the Maximum Mean Discrepancy (MMD). We demonstrate that, for a fixed number of samples, our proposed surrogate exhibits lower variance and thus greater reliability compared to conventional metrics such as Conditional Value at Risk (CVaR). Consequently, the performance of motion planners integrating our surrogate converges faster with increasing data than those employing CVaR.

Basant and Arun are with the University of Tartu. This research was in part supported by grant PSG753 from Estonian Research Council, collaboration project LLTAT21278 with Bolt Technologies and project TEM-TA101 funded by European Union and Estonian Research Council. Emails: basant.sharma@ut.ee, arun.singh@ut.ee, Code: https://github.com/Basant1861/MMD-OPT,https://github.com/Basant1861/MPC-MMD

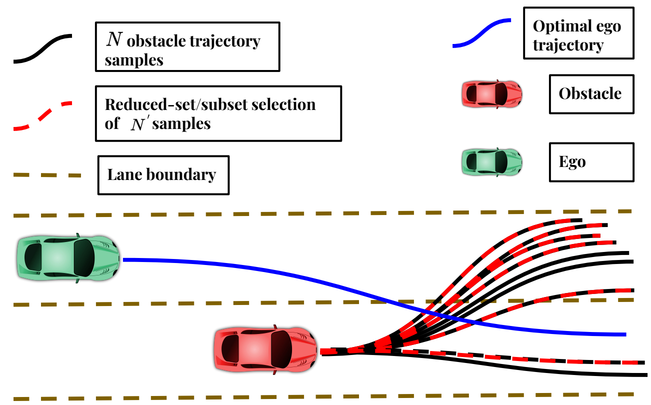

Figure 1: The figure shows a scenario where an obstacle has multiple intents (lane-change vs lane-following), each associated with a trajectory distribution. However, both intents have wildly different probabilities. In this particular example, the probability of lane-change is higher. For safe navigation, 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 a black-box distribution. This allows us to plan probabilistically safe motions while appropriately discriminating the low and high-probability obstacle manoeuvres.

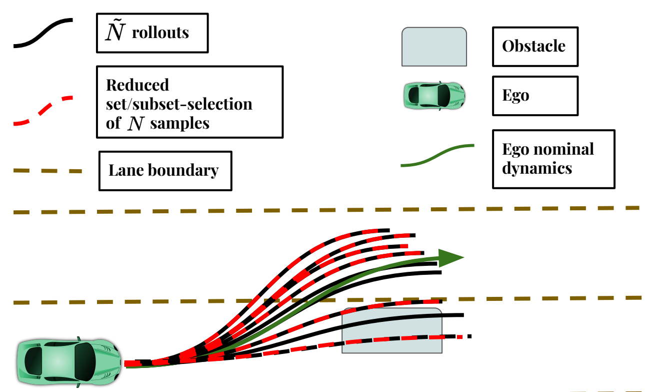

Figure 2: A standard pipeline for risk-aware trajectory optimization for stochastic dynamics based on control sampling along with our improvement. These class of approaches rely on simulating the forward dynamics of the vehicle to obtain $\tilde{N}$ samples from the state trajectory distribution, which are then used to estimate risk. Our work provides a novel risk-surrogate and a systematic way of estimating it using a reduced number (N) of state trajectory samples.

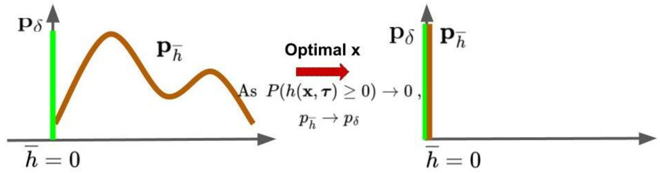

Figure 3: The mass of $p_{\bar{h}}$ is to the right of $\bar{h} = 0$. The optimal control input is one that leads to state-trajectory distribution for which $p_{\bar{h}}$ resembles a Dirac-Delta distribution.

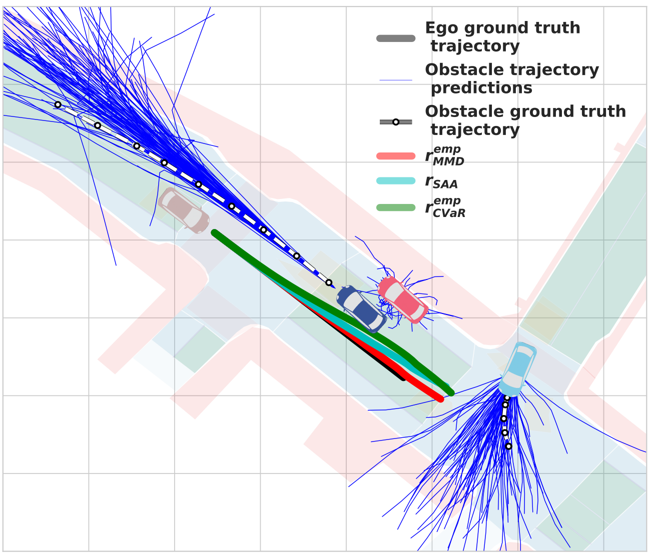

Figure 4: Risk-Aware trajectory planning in an unprotected intersection. The blue trajectories show the trajectory predictions of the neighboring vehicles. As can be seen, the predictions are highly multi-modal capturing different driving intents. Typical ego-vehicle trajectories resulting from different collision risk costs are also shown. It can be seen that r^{emp}_{MMD} 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^{emp}_{CVaR} 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.

II. Problem Formulation

A. Algebraic Form of Risk

Let $h_k(\mathbf{x}_k, \boldsymbol{\tau}_k) \leq 0, \forall k$ denote state-dependent safety constraints, where $\mathbf{x}_k$ is the ego state and $\boldsymbol{\tau}_k$ the obstacle position at time k. Let $\mathbf{x}$ and $\boldsymbol{\tau}$ be obtained by stacking together the respective values at different time step. The worst-case safety constraint value over the horizon is

$$h(\mathbf{x}, \boldsymbol{\tau}) = \max_k \ h_k(\mathbf{x}_k, \boldsymbol{\tau}_k). \tag{1}$$

In the stochastic setting, either $\mathbf{x}$ or $\boldsymbol{\tau}$ (or both) may be random variable, leading to the following general definition of risk:

$$r = P(h(\mathbf{x}, \boldsymbol{\tau}) \geq 0), \tag{2}$$

which measures the probability of violating the safety constraints. The fundamental challenge is that it is generally intractable to obtain a analytical formula for r if the uncertainty in $\mathbf{x}_k$, $\boldsymbol{\tau}_k$ is non-Gaussian. One can always linearize $h(.)$ and compute a Gaussian approximation of the uncertainty, but such approaches are unlikely to provide the reliable risk estimate.

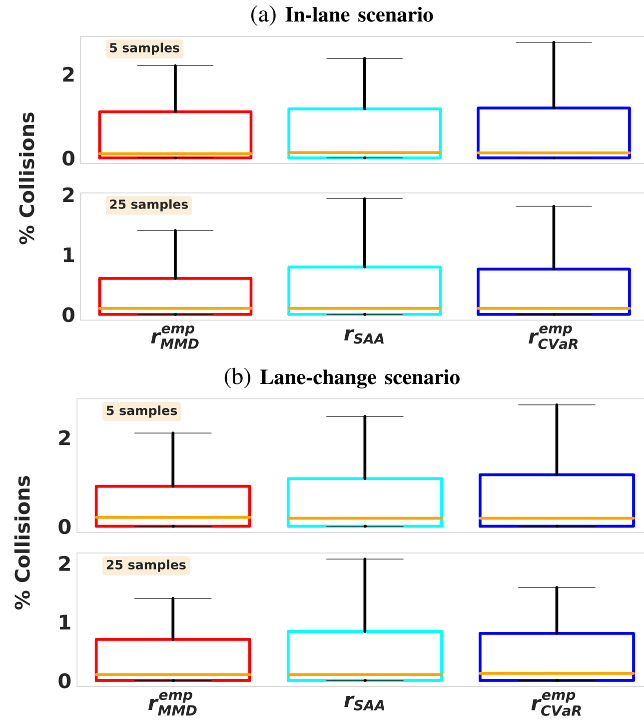

Figure 5: Comparing how well different risks r^{emp}_{MMD}, r_{SAA}, r^{emp}_{CVaR} perform on real-world datasets with Neural Network based trajectory predictors (Trajectron++ [1] in this case). Typically, these predictors define a very complex distribution over the possible trajectories of the neighboring 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^{emp}_{MMD} leads to optimal trajectories with lowest collision-rate in both in-lane driving and lane-change scenarios.

B. MMD based Risk Surrogate

We define the safety constraint residual over the random variable $\mathbf{x}, \boldsymbol{\tau}$ as:

$$\bar{h}(\mathbf{x}, \boldsymbol{\tau}) = \max(0, h(\mathbf{x}, \boldsymbol{\tau})), \tag{3}$$

Let $\bar{h}(\mathbf{x}, \boldsymbol{\tau}) \sim p_{\bar{h}}$. The key insight in our work is that although we don't know the parametric form for $p_{\bar{h}}$, we can be certain that its entire mass lies to the right of $\bar{h} = 0$ (see Fig.3). Moreover, as $P(h(\mathbf{x}, \boldsymbol{\tau}) \geq 0)$ approaches zero, $p_{\bar{h}}$ converges to a Dirac-Delta distribution p_δ. In other words, one way of reducing risk is to minimize the difference between $p_{\bar{h}}$ and p_δ. Intuitively this minimization will make $p_{\bar{h}}$ look similar to p_δ. Thus, we propose the following risk estimate following our prior works [10], [11].

$$r(\mathbf{x}) \approx \mathcal{L}_{dist}(p_\delta, p_{\bar{h}}), \tag{4}$$

where $\mathcal{L}_{dist}$ is any measure that characterizes the difference between two distributions. For example, Kullback-Leibler Divergence (KLD) quantifies distribution similarity but requires known analytical forms, making it unsuitable for comparing $p_{\bar{h}}$ and p_δ using only sample-level information. In the following sections, we propose MMD as a potential choice for $\mathcal{L}_{dist}$.

1) RKHS Embedding of Functions of Random Variables: Our approach builds on the ability of embedding functions of random variables in the RKHS [12], [13]. The RKHS

(a) r^{emp}_{MMD} (b) r^{emp}_{CVaR}

Figure 6: MPC simulations in CARLA. The red line is the route/reference path. Left: trajectory from minimizing r^{emp}_{MMD}. Right: trajectory from minimizing r^{emp}_{CVaR}. With few samples, r^{emp}_{CVaR} varies sharply across MPC steps, often yielding no safe plan (vehicle gets stuck and eventually collides). In contrast, r^{emp}_{MMD} provides a more consistent risk estimate, guiding safe motion.

Table I: $N = 2$, Rollout horizon 40, 50 experiments. Gaussian and Beta noise

Method	Town	% Collisions		% Lane Constr. Viol.		Avg. Speed (m/s)		Max. Speed (m/s)	
		Gaussian	Beta	Gaussian	Beta	Gaussian	Beta	Gaussian	Beta
r^{emp}_{MMD}	T5	**0**	**0**	2.7	3.5	**2.59**	**2.02**	**4.06**	**4.07**
r^{emp}_{CVaR}	T5	15	45	**2.14**	**2.46**	2.08	1.53	3.45	3.4
r^{emp}_{MMD}	T10	**4**	**0**	0.5	5.11	**3.56**	**2**	**7.01**	**3.73**
r^{emp}_{CVaR}	T10	17	4	**0.1**	**3.23**	2.73	1.56	4.55	4.07

embedding of $\overline{h}(\mathbf{x}, \boldsymbol{\tau})$ (or $p_{\overline{h}}$) (function of random variable $(\mathbf{x}, \boldsymbol{\tau})$), and its empirical estimate can be computed as (5a)

$$\mu[\overline{h}] = E[\phi(\overline{h}(\mathbf{x}, \boldsymbol{\tau})] \tag{5a}$$

$$= \int K(\overline{h}(\mathbf{x}, \boldsymbol{\tau}), \cdot) dp_{\overline{h}}(\overline{h}) \tag{5b}$$

where $E[.]$ stands for the expectation operator and ϕ is a non-linear transformation commonly referred to as the feature-map [13]. One of the key properties of ϕ is the so-called kernel trick: the inner product in the RKHS $\left\langle \phi(\mathbf{z}), \phi(\mathbf{z}') \right\rangle_{\mathcal{H}}$ can be expressed as $K_\sigma(\mathbf{z}, \mathbf{z}')$ for any arbitrary vector $\mathbf{z}, \mathbf{z}'$. Here,K_σ is a positive definite function known as the kernel function with hyper-parameter σ. Throughout this paper, we used the Laplacian kernel for which σ represents the kernel-width.

Let δ be a random variable with Dirac-Delta distribution. Let $\mu[\delta]$ be the RKHS embedding of δ (or p_δ). We use the Maximum Mean Discrepancy (MMD) between $p_{\overline{h}}$ and p_δ as our choice for $\mathcal{L}_{dist}$ in (4)

$$r_{MMD} = \mathcal{L}_{dist}(p_\delta, p_{\mathbf{x}}) = \overbrace{\left\|\mu[\overline{h}] - \mu[\delta]\right\|^2_{\mathcal{H}}}^{MMD}. \tag{6}$$

It can be shown that $r_{MMD} = 0$ implies $p_{\overline{h}} = p_\delta$ [12], [13], [14]. In other words, $r_{MMD} = 0$ implies that a state trajectory is safe with probability one. Typically, the r.h.s of (6) is difficult to compute. Thus, it is common to compute the empirical estimate (r^{emp}_{MMD}) through sample mean and utilizing the kernel trick [12], [13], [14], as done in our works [4], [9].

III. Key Results and Conclusion

A. Benchmarking On nuScenes Dataset Using Trajectron++ Predictor [4]

We present results for the case where the ego dynamics are deterministic, and uncertainty stems from neighboring vehicles whose predicted trajectories exhibit multi-modal behaviors We evaluate two driving settings: *in-lane* (ego adjusts only speed) and *lane-change* (full maneuver set). Collision statistics are shown in Fig. 5, with trajectory predictions from Trajectron++ [1]. Fig. 4 illustrates an unprotected intersection where the ego-vehicle must navigate cross-traffic with multi-modal obstacle predictions (e.g., left vs. right turns). The trajectories optimized with r^{emp}_{MMD} align most closely with ground-truth human driving, as r^{emp}_{MMD} correctly prioritizes more likely obstacle maneuvers (e.g., left turn of the cyan vehicle). In contrast, r_{SAA} and r^{emp}_{CVaR} produce conservative deviations due to overemphasis on less likely outcomes.

B. Benchmarking in MPC Setting using CARLA [9]

Next, we present results for the case where obstacle motion is deterministic, but uncertainty stems from process noise in the ego vehicle's dynamics and localization errors. We evaluate the efficacy of r^{emp}_{MMD} in a MPC setting where constant re-planning is done based on the current feedback of ego and the neighbouring vehicle state. Fig.6 presents a typical performance observed with r^{emp}_{MMD} and r^{emp}_{CVaR} under Gaussian noise distribution. Due to extremely small number of constraint residual samples, the variance in r^{emp}_{CVaR} estimation between consecutive MPC steps is large. As a result, the ego-vehicle often gets stuck behind the static obstacles and eventually collides. In sharp contrast, r^{emp}_{MMD} estimates are consistently successful in guiding the ego-vehicle without collision. Table I presents the quantitative benchmarking in the MPC setting. We observed that re-planning can counter some of the effects of noise. As can be seen, r^{emp}_{MMD} outperformed r^{emp}_{CVaR} in collision rate and achieved higher max and average speeds. However, the lane violations of CVaR baseline was lower. However, this was due to shorter runs caused by collisions in several experiments.

References

[1] B. Ivanovic and M. Pavone, "The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2255–2264.

[3] N. Lee, W. Choi, P. Vernaza, C. B. Choy, P. H. Torr, and M. Chandraker, "Desire: Distant future prediction in dynamic scenes with interacting agents," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 336–345.

[4] B. Sharma and A. K. Singh, "Mmd-opt: Maximum mean discrepancy-based sample efficient collision risk minimization for autonomous driving," *IEEE Transactions on Automation Science and Engineering*, vol. 22, pp. 19 051–19 068, 2025.

[5] J. Yin, Z. Zhang, and P. Tsiotras, "Risk-aware model predictive path integral control using conditional value-at-risk," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 7937–7943.

[6] Z. Wang, O. So, K. Lee, and E. A. Theodorou, "Adaptive risk sensitive model predictive control with stochastic search," in *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, ser. Proceedings of Machine Learning Research, vol. 144. PMLR, 07 – 08 June 2021, pp. 510–522. [Online]. Available: https://proceedings.mlr.press/v144/wang21b.html

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

[8] K. N. Tahmasebi, P. Khound, and D. Chen, "A condition-aware stochastic dynamic control strategy for safe automated driving," *IEEE Transactions on Intelligent Vehicles*, 2024.

[9] B. Sharma and A. K. Singh, "Trajectory optimization under stochastic dynamics leveraging maximum mean discrepancy," *IEEE Robotics and Automation Letters*, vol. 10, no. 6, pp. 6079–6086, 2025.

[10] 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 *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 11 087–11 093.

[11] 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 *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2023, pp. 7448–7455.

[12] C.-J. Simon-Gabriel, A. Scibior, I. O. Tolstikhin, and B. Schölkopf, "Consistent kernel mean estimation for functions of random variables," *Advances in Neural Information Processing Systems*, vol. 29, 2016.

[13] B. Schölkopf, K. Muandet, K. Fukumizu, S. Harmeling, and J. Peters, "Computing functions of random variables via reproducing kernel hilbert space representations," *Statistics and Computing*, vol. 25, pp. 755–766, 2015.

[14] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, "A kernel two-sample test," *J. Mach. Learn. Res.*, vol. 13, no. null, p. 723–773, 2012.