

# Stein Variational Guided Model Predictive Path Integral Control: Proposal and Experiments with Fast Maneuvering Vehicles

Kohei Honda<sup>1</sup>, Naoki Akai<sup>1</sup>, Kosuke Suzuki<sup>1</sup>,  
Mizuho Aoki<sup>1</sup>, Hiroataka Hosogaya<sup>1</sup>, Hiroyuki Okuda<sup>1</sup>, and Tatsuya Suzuki<sup>1</sup>

**Abstract**—This paper presents a novel Stochastic Optimal Control (SOC) method based on Model Predictive Path Integral control (MPPI), named Stein Variational Guided MPPI (SVG-MPPI), designed to handle rapidly shifting multimodal optimal action distributions. While MPPI can find a Gaussian-approximated optimal action distribution in closed form, *i.e.*, without iterative solution updates, it struggles with multimodality of the optimal distributions, such as those involving non-convex constraints for obstacle avoidance. This is due to the less representative nature of the Gaussian. To overcome this limitation, our method aims to identify a target mode of the optimal distribution and guide the solution to converge to fit it. In the proposed method, the target mode is roughly estimated using a modified Stein Variational Gradient Descent (SVG-D) method and embedded into the MPPI algorithm to find a closed-form “mode-seeking” solution that covers only the target mode, thus preserving the fast convergence property of MPPI. Our simulation and real-world experimental results demonstrate that SVG-MPPI outperforms both the original MPPI and other state-of-the-art sampling-based SOC algorithms in terms of path-tracking and obstacle-avoidance capabilities.

[https://github.com/kohonda/proj-svg\\_mppi](https://github.com/kohonda/proj-svg_mppi)

## I. INTRODUCTION

Path tracking and obstacle avoidance are essential capabilities required for autonomous mobile robots. These tasks become especially challenging for fast maneuvering vehicles because the optimal action distribution may be multimodal and rapidly shifting. To solve these tasks, sampling-based Model Predictive Control (MPC) [1], [2] is a widely adopted approach that can handle the non-linearity and non-differentiability of the environment, such as system dynamics and cost maps, in contrast to gradient-based methods [3]–[6].

Among various sampling-based MPCs, sampling-based Stochastic Optimal Control (SOC) is a relatively sample-efficient approach that approximates the optimal action distribution as the solution, based on a given prior action distribution [7]–[9]. In particular, Model Predictive Path Integral control (MPPI) [9] stands out as a promising framework because it can estimate a Gaussian-approximated optimal action distribution by analytically minimizing the Kullback-Leibler (KL) divergence. That is, it can find the optimal solution in *closed form* (*i.e.*, without iterative solution updates, similar to gradient descent) when given a sufficient number of samples.

MPPI, however, has limitations in capturing complex optimal distributions due to its less representative of the Gaus-

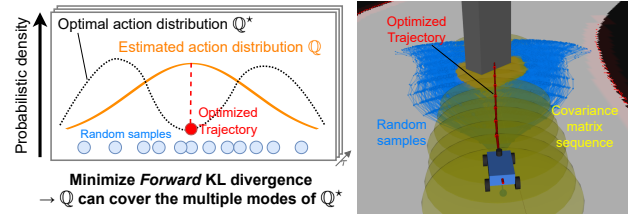


Fig. 1. An open issue of MPPI. The MPPI algorithm minimizes the *Forward* KL divergence instead of the original stochastic optimal control problem. As a result, the estimated action distribution  $\mathbb{Q}$  may cover the multiple modes of the optimal action distribution  $\mathbb{Q}^*$ , and it may lead to finding the collision trajectory as the *optimized* one.

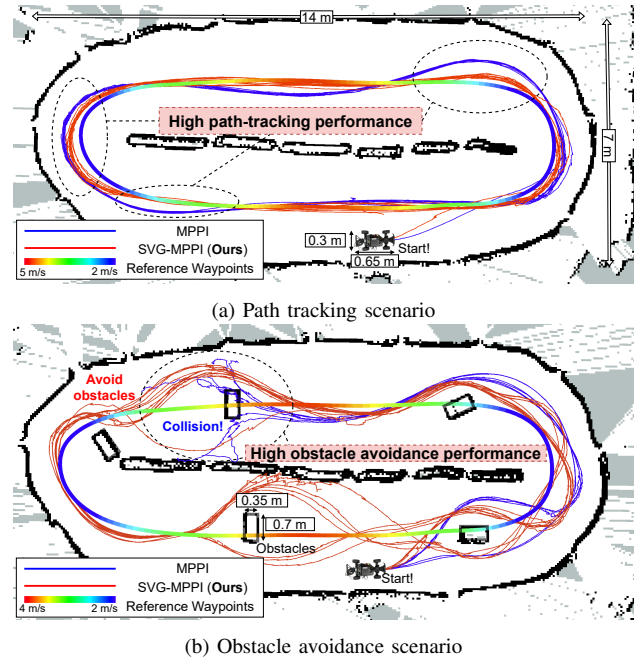


Fig. 2. Vehicle trajectories in the real-world experiment. Our proposed method, named SVG-MPPI, shows smaller curve overshoots than the original MPPI (Fig. 2a) and robustly avoids obstacles where vanilla MPPI collides (Fig. 2b). These results can be found at: <https://www.youtube.com/watch?v=MLaOYQIDL0>

sian. For instance, as illustrated in Fig. 1, an MPPI-based motion planner may produce a collision trajectory as its *optimized* trajectory because the Gaussian-approximated action distribution encompasses the multimodal optimal distribution (see Section III-B for details). To address these shortcomings, several studies have proposed advanced techniques that capture the multimodality using more representative prior distributions [10]–[13] or find *mode-seeking* solutions by leveraging the asymmetry of the KL divergence [14]. However, unlike MPPI, these methods cannot find the so-

\*This work was supported by Tateishi Science and Technology Foundation Research Grants (C) and JSPS KAKENHI Grant Number 23KJ1067.

<sup>1</sup>The Department of Mechanical Systems Engineering, Graduate School of Engineering, Nagoya University, Furo-cho, Chikusa-ku, Nagoya, Aichi, Japan, [honda.kohei.b0@s.mail.nagoya-u.ac.jp](mailto:honda.kohei.b0@s.mail.nagoya-u.ac.jp)

lution in closed form. In other words, they require iterative solution updates with sampling and evaluation to converge the solution. The iterative updates degrade the convergence of the solution compared to MPPI because the solution is suboptimal and the sequential iteration process restricts the parallelism and the number of samples per iteration.

In this paper, we propose Stein Variational Guided MPPI (SVG-MPPI), a sampling-based SOC method based on MPPI to address rapidly shifting multimodal optimal action distributions. Contrary to the existing methods that capture the complex multimodal distribution, our approach aims to narrow down a single target mode within the multimodal distribution and approximate it with a Gaussian distribution by the MPPI algorithm. As a result, our method can obtain a mode-seeking action distribution in closed form, *i.e.*, preserving the fast solution convergence property of MPPI.

Specifically, our proposed method first roughly identifies the target mode by utilizing a small set of samples (hereafter referred to as *guide particles*) and a modified Stein Variational Gradient Descent (SVGd) method [15]. The SVGd method transports the guide particles to near the peak of the target mode within the optimal distribution using its surrogate gradients. Based on the insight that the transport trajectory leading to the peak represents a part of the shape of the target mode, we estimate the rough variance of the target mode using the Gaussian fitting method. Our method then incorporates the peak and variance of the target mode into the MPPI algorithm to converge the solution to cover only the target mode.

In summary, the main contribution of this work is as follows. We propose an MPPI-based SOC method capable of efficiently capturing a single mode of the optimal action distribution. The mode-seeking solution is achieved in closed form by guiding it using a modified SVGd algorithm. Our method has been validated through simulation and real-world experiments, focusing on path tracking and obstacle avoidance tasks for a 1/10th scale vehicle<sup>1</sup>. Empirical results demonstrate that SVG-MPPI outperforms standard MPPI [9] and other state-of-the-art SOC algorithms [13], [14] regarding the path-tracking and obstacle-avoidance performances.

## II. RELATED WORK

To address the rapidly shifting complex optimal action distribution in the sampling-based SOC framework, existing approaches can be broadly categorized into the following three types:

1) *Increasing Effective Samples*: One strategy for improving the MPPI algorithm involves increasing the number of low-state-cost and feasible samples to find better solutions. Some studies propose ancillary controllers that compensate for the MPPI solution [16], [17], and other studies focus on manipulating the prior distribution [18]–[21] to boost the number of the effective samples. However, some approaches

require linearization of the system [16]–[18], and others do not anticipate multimodality [19]–[21].

2) *Variational Inference Approach*: The second approach for capturing the complex optimal action distribution is to use a more representative model, such as Gaussian mixture models, through variational inference methods [10]–[13]. Although these approaches excel at approximating the multimodality of the distribution, they cannot find the solution in closed form and necessitate iterative solution updates, unlike MPPI. Moreover, practical control problems require a deterministic solution for system input rather than a probabilistic multimodal solution. Extracting a smoothly shifting single deterministic solution from the obtained multimodal solution is far from trivial. Therefore, a *mode-seeking* solution is necessary for practical application to control problems.

3) *Finding a Mode-Seeking Solution*: The third approach aims to identify one of the modes of the optimal distribution as the target for convergence, leveraging the asymmetry of the KL divergence [14]. The KL divergence exhibits an asymmetric property when its arguments are reversed. This approach can obtain a mode-seeking solution by directly minimizing KL divergence reversed from the ones used in the original MPPI. However, minimizing the reverse KL divergence cannot be solved in closed form. Our approach can obtain a mode-seeking solution without sacrificing the closed-form optimality inherent in MPPI by guiding the solution with the reverse KL divergence.

## III. MPPI REVIEW

Our aim is to find a mode-seeking action distribution by improving the MPPI algorithm because MPPI struggles with multimodal optimal distributions due to its inherence. In this section, we first review the MPPI theory described in the original literature [9] and then derive the open issue of MPPI to be addressed in this work. Finally, we provide a key property of the KL divergence to address this issue.

### A. Review of the MPPI Theory

1) *Problem Formulation*: The MPPI theory models the control target system as a general discrete time dynamical system  $\mathbf{x}_{\tau+1} = \mathbf{F}(\mathbf{x}_{\tau}, \mathbf{v}_{\tau})$ ,  $\tau \in \{0, \dots, T-1\}$ , where  $\mathbf{x}_{\tau} \in \mathbb{R}^n$  is the predictive state vector at time  $\tau$ ,  $\mathbf{v}_{\tau} \in \mathbb{R}^m$  is the input vector at time  $\tau$ ,  $T$  is the prediction horizon length, and  $\mathbf{F}$  denotes the general non-linear and non-differentiable state transition function. The input vector  $\mathbf{v}_{\tau}$  is injected with Gaussian noise:  $\mathbf{v}_{\tau} \sim \mathcal{N}(\mathbf{u}_{\tau}, \Sigma_{\tau})$ , as a prior action distribution, where  $\mathbf{u}_{\tau} \in \mathbb{R}^m$  is the actual input vector and  $\Sigma_{\tau} = \text{diag}\{\sigma_{\tau}^0, \dots, \sigma_{\tau}^{m-1}\} \in \mathbb{R}^{m \times m}$  is a given covariance matrix of the Gaussian noise.

MPPI aims to estimate an Gaussian action distribution  $\mathbb{Q}$  from randomly generated  $K$  input sample sequences,  $\mathbf{V} = \{\mathbf{V}_k\}_{k=0}^{K-1} \in \mathbb{R}^{K \times T \times m}$ , where  $\mathbf{V}_k = \{\mathbf{v}_{\tau}\}_{\tau=0}^{T-1} \in \mathbb{R}^{T \times m}$ . The Probability Density Function (PDF)  $q(\mathbf{V}_k | \mathbf{U}_t, \Sigma)$  corresponding to  $\mathbb{Q}$  is analytically expressed as:

$$q(\mathbf{V}_k | \cdot) = Z^{-1} \exp \left( -\frac{1}{2} \sum_{\tau=0}^{T-1} (\mathbf{v}_{\tau} - \mathbf{u}_{\tau})^{\top} \Sigma_{\tau}^{-1} (\mathbf{v}_{\tau} - \mathbf{u}_{\tau}) \right), \quad (1)$$

<sup>1</sup>Our system was deployed in a public competition called FITENTH (<https://fitenth.org/>) held at IEEE ICRA 2023, London, where we placed 4th out of 22 teams.

where  $Z = \sqrt{(2\pi)^{mT} |\boldsymbol{\Sigma}|}$  is the normalization term,  $\mathbf{U}_t = \{\mathbf{u}_\tau\}_{\tau=0}^{T-1} \in \mathbb{R}^{T \times m}$  is the actual control input sequence, and  $\boldsymbol{\Sigma} = \text{diag}\{\boldsymbol{\Sigma}_\tau\}_{\tau=0}^{T-1} \in \mathbb{R}^{mT \times mT}$  is a given covariance matrix sequence. To find an optimal control input sequence  $\mathbf{U}_t^*$ , MPPI solves the following SOC problem:

$$\mathbf{U}_t^* = \underset{\mathbf{U}_t \in \mathcal{U}}{\text{argmin}} \mathbb{E}_{\mathbf{V}_k \sim \mathbb{Q}} \left[ \phi(\mathbf{x}_T) + \sum_{\tau=0}^{T-1} \mathcal{L}(\mathbf{x}_\tau, \mathbf{u}_\tau) \right], \quad (2)$$

where  $\mathcal{U}$  is the admissible control set [22],  $\mathbf{x}_0 = \mathbf{x}_t$  is the observed state at time  $t$ ,  $\phi$  is the terminal cost function, and  $\mathcal{L}$  is the stage cost function. Note that the MPPI theory assumes that the stage cost can be divided into state and input costs as,  $\mathcal{L}(\mathbf{x}_\tau, \mathbf{u}_\tau) = c(\mathbf{x}_\tau) + \frac{\lambda}{2} (\mathbf{u}_\tau^\top \boldsymbol{\Sigma}_\tau^{-1} \mathbf{u}_\tau + \boldsymbol{\beta}_t^\top \mathbf{u}_\tau)$ , where  $\boldsymbol{\beta}_t$  is a given parameter independent of  $\mathbf{u}_\tau$ . In the following, the  $k$ th *sequence state cost* is defined as,  $S(\mathbf{V}_k) = \phi(\mathbf{x}_T) + \sum_{\tau=0}^{T-1} c(\mathbf{x}_\tau)$ .

Since (2) is difficult to solve directly, MPPI instead minimizes the *Forward KL* (FKL) divergence  $\mathbb{D}_{\text{KL}}(\mathbb{Q}^* \parallel \mathbb{Q})$  that the action distribution  $\mathbb{Q}$  covers the optimal action distribution  $\mathbb{Q}^*$ , as follows:

$$\mathbf{U}_t^* \simeq \underset{\mathbf{U}_t \in \mathcal{U}}{\text{argmin}} \{\mathbb{D}_{\text{KL}}(\mathbb{Q}^* \parallel \mathbb{Q})\}. \quad (3)$$

2) *The Analytical PDF of the Optimal Action Distribution*: To minimize the FKL divergence in (3), the optimal PDF  $q^*(\mathbf{V}_k)$  of the  $\mathbb{Q}^*$  is derived in the following. First, the *Free-energy* of the system is defined as,  $\mathcal{F}(\mathbf{V}_k) = \log(\mathbb{E}_{\mathbb{P}}[\exp(-\frac{1}{\lambda} S(\mathbf{V}_k))])$ , where  $\lambda$  is a given temperature parameter and  $\mathbb{P}$  is a base distribution, which is roughly analogous to a Bayesian prior. The PDF of the  $\mathbb{P}$  is  $q(\mathbf{V}_k | \tilde{\mathbf{U}}, \boldsymbol{\Sigma})$ , where  $\tilde{\mathbf{U}} = \{\tilde{\mathbf{u}}_\tau\}_{\tau=0}^{T-1} \in \mathbb{R}^{T \times m}$  represents a given nominal control input sequence (hereafter referred to as *nominal sequence*). The Free-energy is then expanded by Jensen's inequality, the definition of FKL divergence, and (1) as,

$$-\lambda \mathcal{F}(\mathbf{V}_k) \leq \mathbb{E}_{\mathbf{V}_k \sim \mathbb{Q}} \left[ \phi(\mathbf{x}_T) + \sum_{\tau=0}^{T-1} \mathcal{L}(\mathbf{x}_\tau, \mathbf{u}_\tau) \right]. \quad (4)$$

Because the right hand side of (4) is equivalent to the cost function in (2),  $-\lambda \mathcal{F}(\mathbf{V}_k)$  represents the lower bound of the original SOC problem. Thus, the optimal PDF  $q^*(\mathbf{V}_k)$  is derived from the equality condition of (4) as,

$$q^*(\mathbf{V}_k | \tilde{\mathbf{U}}, \boldsymbol{\Sigma}) = \eta^{-1} \exp\left(-\frac{1}{\lambda} S(\mathbf{V}_k)\right) q(\mathbf{V}_k | \tilde{\mathbf{U}}, \boldsymbol{\Sigma}), \quad (5)$$

where  $\eta = \int_{\mathbb{R}^{T \times m}} p(\mathbf{V}_k | \lambda) \exp(-\frac{1}{\lambda} S(\mathbf{V}_k)) d\mathbf{V}_k$  is the normalization term.

3) *Forward KL Divergence Mimimization*: As mentioned in section III-A.1, the approach of MPPI is minimizing the FKL divergence in (3), which can be transformed by the definition of the FKL divergence and (1) into the following closed-form expression:

$$\begin{aligned} \mathbf{U}_t^* &= \underset{\mathbf{U}_t \in \mathcal{U}}{\text{argmax}} \mathbb{E}_{\mathbf{V}_k \sim \mathbb{Q}^*} \left[ \frac{1}{2} \sum_{\tau=0}^{T-1} (\mathbf{v}_\tau - \mathbf{u}_\tau)^\top \boldsymbol{\Sigma}_\tau^{-1} (\mathbf{v}_\tau - \mathbf{u}_\tau) \right] \\ &= \mathbb{E}_{\mathbf{V}_k \sim \mathbb{Q}^*} [\mathbf{V}_k]. \end{aligned} \quad (6)$$

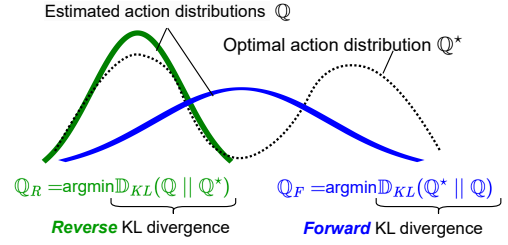


Fig. 3. Asymmetry properties of the KL divergence [14]. The KL divergence has an asymmetric property when its arguments are reversed. Minimizing the *Forward KL* (FKL) divergence results in covering the multiple modes within the optimal action distribution. In contrast, minimizing *Reverse KL* (FKL) divergence converges to a single mode of the optimal distribution, a property known as *mode-seeking*.

Thus, the optimal control input sequence  $\mathbf{U}_t^*$  is obtained by sampling  $\mathbf{V}_k$  from the optimal action distribution  $\mathbb{Q}^*$  and taking its expected value, as in (6). However, since sampling directly from  $\mathbb{Q}^*$  is not possible, importance sampling [23] is used with the  $q^*(\mathbf{V}_k | \tilde{\mathbf{U}}, \boldsymbol{\Sigma})$  in (5), as follows:

$$\begin{aligned} \mathbf{U}_t^* &= \mathbb{E}_{\mathbf{V}_k \sim \mathbb{Q}} \left[ \frac{q^*(\mathbf{V}_k | \tilde{\mathbf{U}}, \boldsymbol{\Sigma})}{q(\mathbf{V}_k | \tilde{\mathbf{U}}, \boldsymbol{\Sigma})} \mathbf{V}_k \right] \simeq \frac{1}{K} \sum_{k=0}^{K-1} w(\mathbf{V}_k) \mathbf{V}_k, \quad (7) \\ w(\mathbf{V}_k) &= \eta^{-1} \exp\left(-\frac{1}{\lambda} S(\mathbf{V}_k) - \sum_{\tau=0}^{T-1} (\hat{\mathbf{u}}_\tau - \tilde{\mathbf{u}}_\tau)^\top \boldsymbol{\Sigma}_\tau^{-1} \mathbf{v}_\tau\right), \quad (8) \end{aligned}$$

where  $\hat{\mathbf{U}} = \{\hat{\mathbf{u}}_\tau\}_{\tau=0}^{T-1}$  is the initial estimated control input sequence, often using the previous optimal solution. Importantly, based on the law of large numbers, the MPPI algorithm can obtain a globally optimal solution in the sense of minimizing the FKL divergence in closed form, *i.e.*, the optimal solution can be obtained by sampling sufficiently large  $K$  samples once and taking their weighted average, without iterative solution updates.

### B. An Open Issue of MPPI

The MPPI algorithm can estimate the Gaussian-approximated optimal distribution by analytically minimizing the FKL divergence in closed form. However, the Gaussian approximation causes an open issue that the estimated action distribution may cover the entire multimodal optimal distribution, as illustrated in Fig. 3 referenced from [14]. For instance, in the context of an obstacle avoidance task, MPPI finds the collision trajectory as an optimal solution, as shown in Fig. 1. This occurs because the two branching paths for avoiding obstacles correspond to the peaks of the two modes within the optimal action distribution, and the estimated action distribution covers both modes. To solve this problem, we need to find a mode-seeking action distribution that covers a single target mode of the multimodal distribution.

### C. Asymmetry Property of the KL divergence

One key to finding the mode-seeking solution is to utilize the asymmetry property of the KL divergence. The standard MPPI algorithm minimizes *Forward KL* (FKL) divergence, denoted as  $\mathbb{D}_{\text{KL}}(\mathbb{Q}^* \parallel \mathbb{Q})$ , where  $\mathbb{Q}$  may cover the entire  $\mathbb{Q}^*$ . In contrast, when the input arguments are reversed, the *Reverse KL* (RKL) divergence  $\mathbb{D}_{\text{KL}}(\mathbb{Q} \parallel \mathbb{Q}^*)$  exhibits a

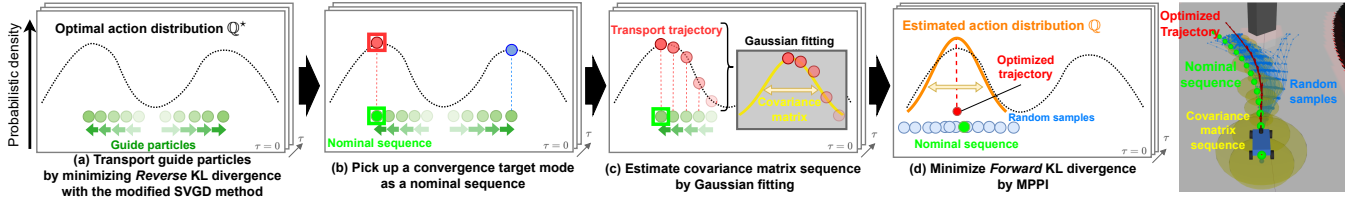


Fig. 4. Overview of the proposed method. (a) SVG-MPPI first transports guide particles by minimizing the RKL divergence using the modified SVGD method. These transport trajectories of the guide particles are employed to roughly estimate a convergent target mode. (b) We then identify the peak of the target mode by simply taking the lowest sequence cost guide particle as a nominal sequence. (c) The variances are also roughly estimated by the Gaussian fitting algorithm with the transport trajectory. (d) Finally, SVG-MPPI minimizes the FKL divergence, incorporating the nominal sequence and adaptive covariance matrices in the process. The SVG-MPPI can efficiently find a mode-seeking action distribution compared to the original MPPI (Fig. 1).

different characteristic, where  $\mathbb{Q}$  converges to a single mode of the optimal distribution  $\mathbb{Q}^*$ , as shown in Fig. 3.

#### IV. STEIN VARIATIONAL GUIDED MPPI

Although some existing methods [13], [14] can find the mode-seeking solution by directly minimizing the RKL divergence described in Section III-C, the solution cannot be obtained in closed form, resulting in a degradation of convergence. To overcome the limitation, our approach guides the MPPI solution to converge to a single target mode roughly identified using a modified SVGD method [15].

As an overview of our method illustrated in Fig. 4, we first spread a small set of samples (guide particles) and transport them by the modified SVGD method. Because the SVGD method iteratively reduces the RKL divergence based on the surrogate gradients of the optimal action distribution  $\mathbb{Q}^*$ , one of the guide particles is transported to near the peak of the target mode. We then roughly estimate the variance of the target mode by fitting the transport trajectory leading to the peak as Gaussian. Finally, we incorporate the peak and variance of the target mode (as the nominal sequence and covariance matrix sequence) into the MPPI algorithm to converge the solution to cover only the target mode.

##### A. Transport Guide Particles by the Modified SVGD method

Our proposed method first spreads  $K_g$  guide particles  $\{\mathbf{V}_k^g\}_{k=0}^{K_g-1} \in \mathbb{R}^{K_g \times T \times m}$  in addition to the  $K$  samples used in the standard MPPI algorithm. Note that  $K_g \ll K$ , in particular, we use  $K_g = 1$  in our experiments in section V. The guide particles are transported to make the RKL divergence smaller by applying the following gradient descent update with  $L$  iterations per control cycle given a small step size  $\varepsilon$ ,

$$\hat{\mathbf{V}}_k^g \leftarrow \mathbf{V}_k^g - \varepsilon \nabla_{\mathbf{V}_k^g} \mathbb{D}_{\text{KL}}(\mathbb{Q} \parallel \mathbb{Q}^*), \quad \forall k \in \{0, \dots, K_g - 1\}. \quad (9)$$

The RKL divergence  $\mathbb{D}_{\text{KL}}(\mathbb{Q} \parallel \mathbb{Q}^*)$  is expanded from its definition and Jensen's inequality as,

$$\begin{aligned} \mathbb{D}_{\text{KL}}(\mathbb{Q} \parallel \mathbb{Q}^*) &= -\mathbb{E}_{\mathbf{V}_k^g \sim \mathbb{Q}} \left[ \log \frac{q^*}{q} \right] = -\mathbb{E}_{\mathbf{V}_k^g \sim \mathbb{Q}} \left[ \log w(\mathbf{V}_k^g) \right] \\ &\leq -\log \mathbb{E}_{\mathbf{V}_k^g \sim \mathbb{Q}} \left[ w(\mathbf{V}_k^g) \right], \end{aligned} \quad (10)$$

where  $w(\cdot)$  is the weight function in (8). Now, we can reduce the RKL divergence of the guide particles by assigning the gradient of (10) into (9) since the right hand side of (10) is the upper bound of the RKL divergence. However, the weight function  $w(\cdot)$  has the sequence state cost  $S(\mathbf{V}_k^g)$ ,

which is generally assumed to be non-differentiable with  $\mathbf{V}_k^g$ . Therefore, we estimate the stochastic surrogate gradient described in [13]:

$$\begin{aligned} \nabla_{\mathbf{V}_k^g} \mathbb{D}_{\text{KL}}(\mathbb{Q} \parallel \mathbb{Q}^*) &= -\nabla_{\mathbf{V}_k^g} \log \mathbb{E}_{\mathbf{V}_k^g \sim \mathbb{Q}} [w(\mathbf{V}_k^g)] \\ &= -\frac{\mathbb{E}_{\mathbf{V}_k^g \sim \mathbb{Q}} \left[ w(\mathbf{V}_k^g) \nabla_{\mathbf{V}_k^g} \log \pi(\mathbf{V}_k^g[i] \mid \mathbf{V}_k^g) \right]}{\mathbb{E}_{\mathbf{V}_k^g \sim \mathbb{Q}} [w(\mathbf{V}_k^g)]} \\ &\simeq -\frac{\sum_{i=0}^{N-1} \left\{ w(\mathbf{V}_k^g[i]) \nabla_{\mathbf{V}_k^g} \log \pi(\mathbf{V}_k^g[i] \mid \mathbf{V}_k^g) \right\}}{\sum_{i=0}^{N-1} w(\mathbf{V}_k^g[i])}, \end{aligned} \quad (11)$$

where  $\nabla x = x \nabla \log x$  is utilized.  $\pi(\mathbf{V}_k^g[i] \mid \mathbf{V}_k^g)$  is the PDF of the distribution for providing the stochastic perturbation, and we use Gaussian distribution in this paper, *i.e.*,  $\mathbf{V}_k^g[i] \sim \mathcal{N}(\mathbf{V}_k^g, \Sigma_g)$ , given a constant covariance matrix  $\Sigma_g$ . The update equation in (9) converges to a simplified version of the update equation using the SVGD method in [13], with the kernel function and log-prior terms removed. These removed terms penalize particle cohesion and preserve the diversity of the samples. However, for our approach, the penalty terms are arbitrarily removed because we expect the guide particles to aggregate to some of the peaks of the optimal distribution. Note that the computational complexity for the  $L$  times transport in (9) is  $\mathcal{O}(LK_g N)$  and relatively tiny compared to the MPPI's one (*i.e.*,  $\mathcal{O}(TKm)$ ) when  $K_g \ll K$ .

##### B. FKL divergence Minimization with the Nominal Sequence and Adaptive Covariance Matrix Sequence

We now obtain  $K_g$  transport trajectories for each guide sample as,  $\tau_k^g = \{\mathbf{V}_k^g[l]\}_{l=0}^{L-1} \in \mathbb{R}^{L \times T \times m}$ ,  $k \in \{0, \dots, K_g - 1\}$ , where  $L$  is the number of iterations for (9). We then roughly identify the peak and variance of the target mode using the transport trajectories and incorporate them as the nominal sequence  $\tilde{\mathbf{U}}$  and covariance matrix sequence  $\tilde{\Sigma}$  in the MPPI algorithm described in Section III-A. Importantly, our method is naturally incorporated into the MPPI framework without changing its theory, so it retains the advantages of MPPI; its solution is obtained in closed form.

1) *Picking a Target Mode:* The guide particles transported in sufficiently large iterations  $L$  are expected to be close to some of the peaks. Since the convergence goal of a gradient descent depends on the initial value, the guide particles can be transported to different peaks within the optimal distribution after the iterations. The strategy for determining the peak of the target mode is simply to pick

up a transported guide particle that has the lowest sequence state cost. Therefore, the peak, as the nominal sequence  $\tilde{U}$  of the MPPI algorithm, is determined, as follows:

$$\tilde{U} = \mathbf{V}_{k^*}^g, \text{ where } k^* = \underset{k}{\operatorname{argmin}} \{S(\mathbf{V}_k^g[L-1])\}_{k=0}^{K_g-1}. \quad (12)$$

The nominal sequence  $\tilde{U}$  is used in the weight function (8), which penalizes the deviation of the nominal sequence (the peak of the target mode) and the mean of the action distribution, which corresponds to determining the coverage target mode within the optimal action distribution.

2) *Estimation of the Adaptive Covariance Matrix Sequence*: To converge the action distribution to cover only the target mode, we should set an appropriate covariance matrix sequence  $\Sigma$ , which is used in the prior action distribution of random samples and weight function in (8) for the MPPI algorithm. Our method utilizes the entire  $k^*$ th transport trajectory  $\tau_{k^*}^g$  to estimate the rough variance of the target mode because the transport trajectory leading to the peak is expected to represent a part of the shape of the target mode. Specifically, we use Gao's algorithm [24], which is known as a fast and rough Gaussian fitting method. For all predictive times and control inputs  $\forall \tau \in [0, \dots, T-1]$  and  $\forall i \in [0, \dots, m-1]$ , let the input vector be  $\mathbf{a} = [\mathbf{V}_{k^*}^g[0, \tau, i], \dots, \mathbf{V}_{k^*}^g[L-1, \tau, i]] \in \mathbb{R}^L$  and the corresponding  $q^*$  in (5) be  $\mathbf{b} = [q^*(\mathbf{V}_{k^*}^g[0]), \dots, q^*(\mathbf{V}_{k^*}^g[L-1])] \in \mathbb{R}^L$ , and the variance  $\Sigma_\tau[i] = \sigma_\tau^i$  for each control input can be estimated by solving the first-order equation in the Gao's algorithm,

$$\sigma_\tau^i = \sqrt{\frac{-1}{2z_2}}, \text{ where} \quad (13)$$

$$\begin{bmatrix} \Sigma \mathbf{b}^2 & \Sigma \mathbf{a} \mathbf{b}^2 & \Sigma \mathbf{a}^2 \mathbf{b}^2 \\ \Sigma \mathbf{a} \mathbf{b}^2 & \Sigma \mathbf{a}^2 \mathbf{b}^2 & \Sigma \mathbf{a}^3 \mathbf{b}^2 \\ \Sigma \mathbf{a}^2 \mathbf{b}^2 & \Sigma \mathbf{a}^3 \mathbf{b}^2 & \Sigma \mathbf{a}^4 \mathbf{b}^2 \end{bmatrix} \begin{bmatrix} z_0 \\ z_1 \\ z_2 \end{bmatrix} = \begin{bmatrix} \Sigma \mathbf{b}^2 \log \mathbf{b} \\ \Sigma \mathbf{a} \mathbf{b}^2 \log \mathbf{b} \\ \Sigma \mathbf{a}^2 \mathbf{b}^2 \log \mathbf{b} \end{bmatrix}.$$

As a result, when the change of the  $q^*$  due to the transport is small, the variance is estimated to be small. In this case, the sequence state cost of the guide particle is largely shifted by the transport. In the opposite case, the variance is kept to be large, as shown in Fig. 4.

## V. EXPERIMENTS

### A. Experiment Setup

We implemented our method for a path-tracking and obstacle-avoidance motion planner of a 1/10th-scale vehicle. The motion planner controls the vehicle's steering angle based on given reference waypoints (path and speed profile), the vehicle's position and velocity, and a cost map constructed from onboard 2D LiDAR scans for obstacles and course detection.

1) *MPC Formulation*: The implemented vehicle dynamics prediction model  $F$  was the kinematic bicycle model [25] considering the dead time and first-order delays with the steering angle input, which have non-linearity. The non-differentiable sequence state cost was  $S(\mathbf{V}_k) = \sum_{\tau=0}^{T-1} \{\Delta d_r(\tau)^2 + 0.01 \Delta \theta_r(\tau)^2 + 1000 \mathbf{1}^{\text{collide}}(\mathbf{x}_\tau)\}$ , where  $\Delta d_r$  and  $\Delta \theta_r$  were position and yaw angle deviations from the reference waypoints, and  $\mathbf{1}^{\text{collide}}$  is a binary indicator function

for collision, provided by the cost map. Please refer to the source code<sup>2</sup> for the detailed algorithm and settings.

2) *Validation Scenarios*: To validate the performance of the proposed method, we conducted tests under two types of driving scenarios, both in simulation and in the real world:

- Path Tracking (PT): The vehicle completed  $N_{\text{pt}}$  laps of the course without encountering unforeseen obstacles.
- Obstacle Avoidance (OA): The vehicle drove  $N_{\text{oa}}$  laps of the course, encountering five unforeseen static obstacles. These obstacles were randomly repositioned around reference waypoints on each lap.

3) *Evaluation Metrics*: To evaluate the path-tracking and obstacle-avoidance capabilities, we employ the following two evaluation metrics:

- Mean Sequence state cost  $S(U_t)$  per lap, denoted as MS: A lower value of this metric indicates better performance in path tracking and obstacle avoidance. In scenario PT, this metric is primarily focused on evaluating path-tracking performance. Conversely, scenario OA, mainly evaluates the obstacle-avoidance performance.
- Collision Rate (CR): This represents the rate of collisions with obstacles or the course, relative to the total number of obstacles encountered in scenario OA.

4) *Baseline Methods*: We implemented the vanilla MPPI [9], Reverse-MPPI with sample rejection [14], and SV-MPC [13] to compare their performance to our proposed method, SVG-MPPI. Reverse-MPPI minimizes the RKL divergence based on mirror descent algorithm [26], and SV-MPC also directly minimizes RKL divergence with variational inference approach, *i.e.*, using SVGD method [15]. Unlike the MPPI that minimizes the FKL divergence, they cannot be solved in closed form, and the number of random samples for the iterative updates must be set smaller than the MPPI and SVG-MPPI to keep real-time processing.

All algorithms, including SVG-MPPI, were developed in ROS and C++ and employed CPU multi-threading via OpenMP [27]. The same MPC formulation and common parameters were used across all algorithms, except the number of random samples. To keep a calculation time of approximately 20 ms on a desktop PC equipped with an Intel Core i9-10850K CPU, we set the number of random samples for MPPI, Reverse-MPPI, SV-MPC, and SVG-MPPI at 10k, 200, 500, and 8k, respectively.

### B. Simulation Results

We executed each method for 100 laps in both the PT and OA scenarios to provide a comprehensive comparison ( $N_{\text{pt}} = N_{\text{oa}} = 100$ ). For the simulation environment, we customized the F1TENTH Gym environment [28] to include unforeseen five obstacles that were randomly replaced around reference waypoints (within 0.1 m) on each lap.

1) *Comparison with Baseline Methods*: Figure 5 shows the evaluation results for the MS metrics in both PT and OA scenarios, and Table I lists the CRs for each method in the OA scenario. In the case of vanilla MPPI, a lower

<sup>2</sup>[https://github.com/kohonda/proj-svg\\_mppi](https://github.com/kohonda/proj-svg_mppi)

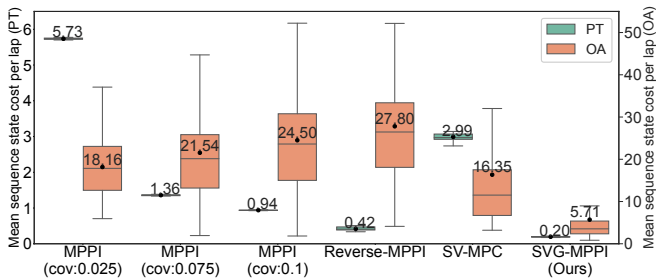


Fig. 5. Comparison of the mean sequence state cost per lap (MS) in the simulation.

TABLE I  
COLLISION RATE (CR) FOR EACH METHOD IN THE SIMULATION

Method	MPPI	MPPI	MPPI	Reverse	SV	SVG-MPPI
Cov. [rad]	0.025	0.075	0.1	MPPI	MPC	(Ours)
CR [%]	13.6	21.0	24.4	24.8	12.4	<b>4.0</b>

covariance of the control input correlates with lower MSs and CRs in the OA scenario; however, it negatively impacts MSs in PT. These findings suggest that vanilla MPPI faces a trade-off between PT and OA capability depending on the fixed covariance parameter. This is attributable to the fact that a small covariance prevents the solution from covering the entire optimal action distribution in the OA, but is not suitable for PT, which requires covering the entire single-mode distribution to track sharp curves rapidly.

Unlike vanilla MPPI, both Reverse-MPPI and SV-MPC aim to minimize RKL divergence directly. As shown in Fig. 5, Reverse-MPPI outperforms vanilla MPPI in PT, while SV-MPC outperforms in the OA scenario. However, each method underperforms in the other scenario; Reverse-MPPI in OA and SV-MPC in PT. In other words, both algorithms have the potential to outperform vanilla MPPI in specific tasks but also exhibit the trade-off between PT and OA capabilities. Note that these methods adaptively estimate the covariance matrices in their algorithms.

On the other hand, our proposed method, SVG-MPPI, demonstrates superior performance in both scenarios, with average MS values of 0.20 and 5.71 for PT and OA, respectively as shown in Fig. 5. These results confirm that SVG-MPPI excels in both PT and OA tasks, avoiding the trade-offs observed in all baseline methods. Furthermore, SVG-MPPI records the lowest CR of 4.0 % across all trials in the OA scenario in Table I. In the rare cases where collisions occurred, SVG-MPPI did not find a mode-seeking solution due to its inability to guarantee one.

2) *Ablation Study*: SVG-MPPI is an extension of vanilla MPPI that incorporates both a Nominal Sequence (NS) and Adaptive Covariance matrix sequence estimation (AC). To demonstrate the complementary nature of these two extensions, we conducted an ablation study for the PT and OA scenarios in the simulation. The results are presented in Table II. Both MPPI and MPPI+NS employed a fixed steering covariance of 0.75 rad. As indicated in Table II, MPPI+NS outperforms vanilla MPPI in both PT and OA scenarios. Moreover, MPPI+AC exhibits even better

TABLE II  
ABLATION STUDY: AVERAGE OF MSS FOR 100 TRIALS

method	Scenario PT	Scenario OA
MPPI (baseline)	1.36	21.54
MPPI + nominal sequence (NS)	<b>0.18</b>	11.03
MPPI + adaptive covariance (AC)	18.18	7.08
MPPI + NS + AC (SVG-MPPI)	0.20	<b>5.71</b>

performance in OA compared to MPPI+NS, although its performance in PT is significantly diminished. When both enhancements (MPPI+NS+AC) are incorporated, SVG-MPPI delivers the best performance in OA and ranks second-best in PT. Overall, these results confirm that the NS and AC functions are complementary and contribute to the enhanced path-tracking and obstacle-avoidance performance.

### C. Real-world Experiment Results

We also conducted real-world experiments using a 1/10th-scale vehicle equipped with a 2D LiDAR (HOKUYO UTM30-LX) and a compact computer (Intel NUC with a Core i7-1360P Processor). To estimate the vehicle's pose on the pre-built map, we employed an advanced Monte Carlo localization method [29].

Figure 2 shows the trajectories of both vanilla MPPI (steering covariance: 0.075 rad) and SVG-MPPI in PT and OA scenarios with a fixed obstacle layout. The figure reveals that SVG-MPPI exhibits smaller overshoots on curves compared to vanilla MPPI, as further illustrated in Fig. 2a. Additionally, Fig.2b indicates that SVG-MPPI successfully avoids obstacles where vanilla MPPI encounters collisions.

For more evaluation, we completed 10 laps in scenario PT ( $N_{pt} = 10$ ) and 3 laps for each of six different static obstacle layouts in scenario OA ( $N_{pt} = 18$ ). In scenario PT, the mean MS values were recorded at 5.83 and 2.03 for vanilla MPPI and SVG-MPPI, respectively. Turning our attention to scenario OA, the mean MS and CR values were (62.73, 55.6 %) for vanilla MPPI and (27.82, 15.4 %) for SVG-MPPI. Overall, SVG-MPPI demonstrated superior path-tracking and obstacle-avoidance capabilities compared to the vanilla MPPI method in real-world testing.

## VI. CONCLUSION AND POTENTIAL LIMITATION

This paper presents a novel MPPI-based SOC method to address the rapidly shifting multimodal optimal action distribution. Our method is capable of finding a mode-seeking solution in closed form by guiding the convergence target of the MPPI solution using the SVGD method. Our simulation and real-world experimental results demonstrate that our method achieves superior path-tracking and obstacle-avoidance capabilities compared to other baseline methods, without the trade-offs of these capabilities commonly observed. While our experiments do not show any performance degradation, a potential limitation of our method arises in cases where the gradients of the optimal distribution are zero outside the peaks of the modes. This limitation is due to the fact that the SVGD method primarily tracks gradients, and the solution may become trapped at such terrace-like places within the optimal distribution.

## REFERENCES

- [1] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robotics & Automation Magazine*, vol. 4, no. 1, pp. 23–33, 1997.
- [2] T. M. Howard, C. J. Green, A. Kelly, and D. Ferguson, "State space sampling of feasible motions for high-performance mobile robot navigation in complex environments," *Journal of Field Robotics*, vol. 25, no. 6-7, pp. 325–345, 2008.
- [3] M. Diehl and S. Gros, "Numerical optimal control," *Optimization in Engineering Center*, 2011.
- [4] J. A. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi: a software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, pp. 1–36, 2019.
- [5] T. Ohtsuka, "A continuation/gmres method for fast computation of nonlinear receding horizon control," *Automatica*, vol. 40, no. 4, pp. 563–574, 2004.
- [6] K. Honda, H. Okuda, T. Suzuki, and A. Ito, "MPC Builder for autonomous drive: Automatic generation of MPCs for motion planning and control," in *Intelligent Vehicles Symposium*. IEEE, 2023, pp. 1–8.
- [7] N. Hansen, S. D. Müller, and P. Koumoutsakos, "Reducing the time complexity of the derandomized evolution strategy with covariance matrix adaptation (CMA-ES)," *Evolutionary computation*, vol. 11, no. 1, pp. 1–18, 2003.
- [8] Z. I. Botev, D. P. Kroese, R. Y. Rubinstein, and P. L' Ecuyer, "The cross-entropy method for optimization," in *Handbook of statistics*. Elsevier, 2013, vol. 31, pp. 35–59.
- [9] G. Williams, P. Drews, B. Goldfain, J. M. Rehg, and E. A. Theodorou, "Information-theoretic model predictive control: Theory and applications to autonomous driving," *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1603–1622, 2018.
- [10] M. Okada and T. Taniguchi, "Variational inference MPC for Bayesian model-based reinforcement learning," in *Conference on robot learning*. PMLR, 2020, pp. 258–272.
- [11] Z. Wang, O. So, J. Gibson, B. Vlahov, M. S. Gandhi, G.-H. Liu, and E. A. Theodorou, "Variational inference MPC using tsallis divergence," *Robotics: Science and Systems*, 2021 (in print).
- [12] L. Barcelos, A. Lambert, R. Oliveira, P. Borges, B. Boots, and F. Ramos, "Dual online Stein variational inference for control and dynamics," *Robotics: Science and Systems*, 2021 (in print).
- [13] A. Lambert, F. Ramos, B. Boots, D. Fox, and A. Fishman, "Stein variational model predictive control," in *Conference on Robot Learning*. PMLR, 2021, pp. 1278–1297.
- [14] T. Kobayashi and K. Fukumoto, "Real-time sampling-based model predictive control based on reverse kullback-leibler divergence and its adaptive acceleration," *arXiv preprint arXiv:2212.04298*, 2022.
- [15] Q. Liu and D. Wang, "Stein variational gradient descent: A general purpose bayesian inference algorithm," *Advances in Neural Information Processing Systems*, vol. 29, 2016.
- [16] G. Williams, B. Goldfain, P. Drews, K. Saigol, J. M. Rehg, and E. A. Theodorou, "Robust sampling based model predictive control with sparse objective information," in *Robotics: Science and Systems*, vol. 14, 2018, p. 2018.
- [17] I. M. Balci, E. Bakolas, B. Vlahov, and E. A. Theodorou, "Constrained covariance steering based Tube-MPPI," in *American Control Conference*. IEEE, 2022, pp. 4197–4202.
- [18] J. Yin, Z. Zhang, E. Theodorou, and P. Tsiotras, "Trajectory distribution control for model predictive path integral control using covariance steering," in *International Conference on Robotics and Automation*. IEEE, 2022, pp. 1478–1484.
- [19] J. Sacks and B. Boots, "Learning sampling distributions for model predictive control," in *Conference on Robot Learning*. PMLR, 2023, pp. 1733–1742.
- [20] D. M. Asmar, R. Senanayake, S. Manuel, and M. J. Kochenderfer, "Model predictive optimized path integral strategies," in *International Conference on Robotics and Automation*. IEEE, 2023, pp. 3182–3188.
- [21] I. S. Mohamed, J. Xu, G. Sukhatme, and L. Liu, "Towards efficient MPPI trajectory generation with unscented guidance: U-MPPI control strategy," *arXiv preprint arXiv:2306.12369*, 2023.
- [22] E. C. Kerrigan, "Robust constraint satisfaction: Invariant sets and predictive control," Ph.D. dissertation, University of Cambridge UK, 2001.
- [23] T. Kloek and H. K. Van Dijk, "Bayesian estimates of equation system parameters: an application of integration by monte carlo," *Econometrica: Journal of the Econometric Society*, pp. 1–19, 1978.
- [24] H. Guo, "A simple algorithm for fitting a gaussian function," *IEEE Signal Processing Magazine*, vol. 28, no. 5, pp. 134–137, 2011.
- [25] R. Rajamani, *Vehicle dynamics and control*. Springer Science & Business Media, 2011.
- [26] A. Beck and M. Teboulle, "Mirror descent and nonlinear projected subgradient methods for convex optimization," *Operations Research Letters*, vol. 31, no. 3, pp. 167–175, 2003.
- [27] R. Chandra, *Parallel programming in OpenMP*. Morgan kaufmann, 2001.
- [28] M. O'Kelly, H. Zheng, D. Karthik, and R. Mangharam, "F1tenth: An open-source evaluation environment for continuous control and reinforcement learning," *Proceedings of Machine Learning Research*, vol. 123, 2020.
- [29] N. Akai, "Reliable monte carlo localization for mobile robots," *Journal of Field Robotics*, vol. 40, no. 3, pp. 595–613, 2023.