

Bridging Active Exploration and Uncertainty-Aware Deployment Using Probabilistic Ensemble Neural Network Dynamics

Taekyung Kim*, Jungwi Mun*, Junwon Seo, Beomsu Kim, and Seongil Hong
 AI Autonomy Technology Center, Agency for Defense Development

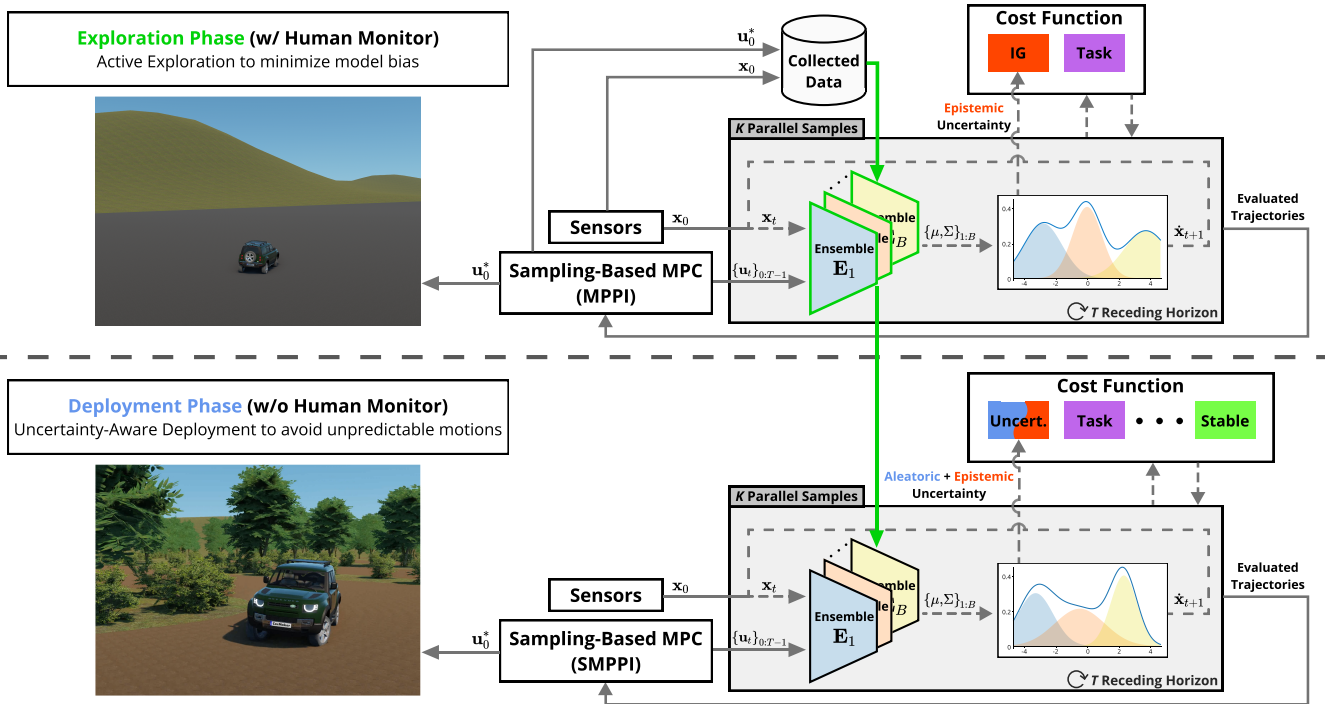


Fig. 1: Overview diagram of our unified model-based reinforcement learning framework with dynamics learning. In *exploration phase*, a parallelized ensemble neural network serves as the robot dynamics and outputs the estimated posterior distribution of the next state. To enable active exploration, we quantify epistemic uncertainty by measuring the ensemble disagreement via Jensen-Rényi Divergence. In *deployment phase*, the neural network dynamics trained during the active exploration phase is applied directly to perform uncertainty-aware control. We transfer the neural network dynamics for uncertainty-aware deployment with minimal modification. Project page: <https://taekyung.me/rss2023-bridging>

Abstract—In recent years, learning-based control in robotics has gained significant attention due to its capability to address complex tasks in real-world environments. With the advances in machine learning algorithms and computational capabilities, this approach is becoming increasingly important for solving challenging control problems in robotics by learning unknown or partially known robot dynamics. Active exploration, in which a robot directs itself to states that yield the highest information gain, is essential for efficient data collection and minimizing human supervision. Similarly, uncertainty-aware deployment has been a growing concern in robotic control, as uncertain actions informed by the learned model can lead to unstable motions or failure. However, active exploration and uncertainty-aware deployment have been studied independently, and there is limited

literature that seamlessly integrates them. This paper presents a unified model-based reinforcement learning framework that bridges these two tasks in the robotics control domain. Our framework uses a probabilistic ensemble neural network for dynamics learning, allowing the quantification of epistemic uncertainty via Jensen-Rényi Divergence. The two opposing tasks of exploration and deployment are optimized through state-of-the-art sampling-based MPC, resulting in efficient collection of training data and successful avoidance of uncertain state-action spaces. We conduct experiments on both autonomous vehicles and wheeled robots, showing promising results for both exploration and deployment.

I. INTRODUCTION

Learning-based robotic control has been an increasingly active research topic in recent decades, pushing the boundary

*These authors contributed equally to this work

of robotic control tasks with new capabilities through model-free reinforcement learning [1, 2, 3, 4], model-based reinforcement learning [5, 6, 7, 8, 9, 10, 11], and robot dynamics learning [12, 13, 14, 15]. There are two key aspects that are commonly taken into account while deploying an intelligent robot. The first is to assure safety, and the second, if at all possible, is to achieve maximum performance.

Model-based methods have been regarded as being advantageous in robotic applications over pure model-free strategies for various reasons. Firstly, sample efficiency is essential since real-world samples are highly expensive in terms of time, labor, and finances [16]. Secondly, humans are primarily more intuitive about how to incorporate prior knowledge into a model compared to a policy or value function [14, 17]. Lastly, models are task-agnostic and may thus be utilized to optimize arbitrary cost functions, whereas the majority of model-free policies are bounded to a specific task.

These benefits have led to recent control paradigms that include dynamics learning into a model-based control framework, outperforming the conventional methods that rely on manually tuned models based on human insights into physics [10, 18, 19]. This scheme may also be described as a model-based reinforcement learning problem, as the model is continuously trained to provide a more accurate representation of the system, while an optimization-based controller serves as the policy. Such a control framework possesses the inherent flexibility to dynamically adjust during execution, for instance, to changing target speeds or actuator limits. Gaussian Process (GP) is a typical regression method in model learning problems due to its advantageous features, such as the ability to cope with small data samples and incorporate prior knowledge [6, 20]. However, it does not scale well to large datasets and to high-dimensional systems. To mitigate this issue, recent robotics studies have benefited from neural networks owing to their universal function approximation properties with high model capacity.

The limitation of these model-based control frameworks is that the accuracy of the learned model drastically affects the control performance [21]. However, the real-world is fraught with uncertainty, making the model learning problem complex and challenging. There are two primary viewpoints within the robotics community that tackle inevitable uncertainty in opposing ways.

A spontaneous solution would be to prevent the robot from entering uncertain state-action spaces to evade unpredictable motions. This is a straightforward strategy, given that an inaccurate prediction of the dynamics model might result in substantial losses in control performance and, in the worst case, instability. We refer to such an uncertainty-averse strategy as *uncertainty-aware deployment* (or safety-aware deployment), and it is usually applied during execution after completing the robot’s training.

The exact opposite of the above strategy is to deliberately visit unexplored state-action spaces that provide high uncertainty. The hypothesis is that by trying to explore uncertain segments of the learned model, it can learn more precise

representation of the system with fewer samples and discover behaviors with higher performance. Such uncertainty-seeking strategy is known as *active exploration*, and it is considered during the early phase of training data exploration.

We argue that these opposite yet analogous strategies addressing uncertainty must be entailed concurrently in a learning-based control framework. No matter how much data is gathered through active exploration, the uncertainty that eventually has not been eliminated needs to be monitored and avoided throughout robot deployment. Conversely, when we take the model uncertainties for risk-averse planning, they are only meaningful if the model is trained with the right data; otherwise, the controller will produce erroneous motions. Nevertheless, these two approaches have largely been studied independently, and there is little literature that integrates them in a seamless manner.

This paper introduces a unified model-based reinforcement learning framework in a robot dynamics learning setting, seamlessly bridging active exploration and uncertainty-aware deployment to fulfill both safety and maximum performance. Fig. 1 shows that the proposed framework consists of two phases. In *exploration phase*, a parallelized ensemble neural network serves as the robot dynamics and outputs the estimated posterior distribution of the next state. In *deployment phase*, the neural network dynamics trained during the active exploration phase is applied directly to perform uncertainty-aware control. Both tasks are optimized using the state-of-the-art sampling-based Model Predictive Control (MPC), which, owing to its property, allows the insertion of arbitrary cost functions after training.

In summary, the contributions of this paper are as follows:

- We introduce a fully parallelized probabilistic ensemble neural network that is sensitive to uncertainty in order to learn robot dynamics.
- We separate epistemic uncertainty from model disagreement for active exploration.
- We transfer the neural network dynamics for uncertainty-aware deployment with minimal modification.
- We conduct extensive experiments with our framework on both autonomous vehicles and wheeled robots, outperforming the compared methods by a large margin.

II. RELATED WORKS

A. Model Learning with Uncertainty

A regression model capable of accurately capturing the nonlinear dynamics is essential for model-based control in uncertain environments with nonlinearity. Data-driven models approximating the robot’s forward dynamics have been extensively studied in numerous prior research, including GP [6, 20], time-varying linear models [22], set-membership identification [23, 24], locally weighted projection regression [25, 26], etc. GP has long been favored in dynamics learning given its simplicity and its inherent nature to account for both types of uncertainty: aleatoric uncertainty and epistemic uncertainty.

Aleatoric uncertainty arises from intrinsic stochasticities present in a system, such as observation noise and process

noise. Epistemic uncertainty, however, originates from a lack of sufficient data to uniquely determine the underlying characteristics of the target system in an approximated model. An ideal model trained with an infinite number of data may have zero epistemic uncertainty, but it is not feasible to eradicate it completely in a realistic setting.

Neural networks are known for being tractable to large training datasets while assuring constant-time inference and having the potential to learn highly complex and even non-smooth robot dynamics. The recent breakthrough in the field of dynamics learning is that neural networks may also capture both types of uncertainty by using an ensemble of probabilistic neural networks [27]. Vlastelica et al. [28] proposed to quantify epistemic uncertainty by taking empirical variances over the means and variances of the propagated posterior distribution of each ensemble. We investigate further to determine a more exact epistemic uncertainty in order to provide a better exploration bonus (which is discussed in Section IV-A).

B. Active Exploration

Human-supervised data collection is time-consuming and can be biased by the fact that manual control of the robot may not encompass all of its possible motions. As a remedy, it was suggested that the agent should plan exploration by directing itself to the states that yield the largest information gain, therefore performing active exploration [29]. In general, the measured model uncertainty is leveraged as the exploration motivation to minimize the required amount of data to gather [30].

We argue that the exploration strategy should be *imaginary* and *online*. If the agent is *reactive* to immediate curiosity [31], which is the opposite to being *imaginary*, it disregards the long-term effects of control actions to the system. Consequently, it must struggle in realms of underactuated continuous control [32]. In addition, since most robotic systems are intrinsically stochastic, the real trajectories would quickly diverge from the planned ones. Therefore, it is desirable to plan future trajectory *online* at every time step, rather than relying on the optimized trajectory found *offline* as in [32].

Several studies address model-based active exploration based on the above hypothesis. Examples of this line of research include the work in [33, 34, 35], in which model-free policies are employed to maximize expected information gain in a specific task. In [36], GP is used for model learning, while a gradient-based MPC handles the information gain acquired from the GP model. In this work, we adopt sampling-based MPC as a policy and demonstrate that it is scalable to handle uncertainty derived from neural networks.

C. Safe Deployment

Uncertainty-averse deployment has been addressed as a major concern throughout robotic control history, and there is a growing interest in its application to model-based reinforcement learning. The uncertain actions informed by the learned model may result in unstable motions or even catastrophic failure.

Kahn et al. [9] incorporated the estimated uncertainty as a cost to learn collision avoidance. Yu et al. [37] penalized uncertainty of the dynamics in an offline reinforcement learning setting to resolve the distributional shift issue. Wang et al. [38] treated trajectories having greater uncertainty than a certain threshold as constraint conditions in an application of rough terrain navigation. Similar to [38], we penalize uncertainty with both soft and hard costs to maximize the control performance while ensuring safety.

III. METHODS

A. Probabilistic Ensemble Neural Network

Consider a general continuous time system dynamics $\dot{\mathbf{x}}_{t+1} = \mathbf{F}(\mathbf{x}_t, \mathbf{u}_t)$ where \mathbf{F} is a nonlinear function, $\mathbf{x}_t \in \mathbb{R}^n$ and $\mathbf{u}_t \in \mathbb{R}^m$ are the observed state vector and applied action input at time t , respectively. To account for both aleatoric and epistemic uncertainty, we use Probabilistic Ensemble Neural Network (PENN) to approximate the unknown dynamics \mathbf{F} via data-driven manner [27, 39].

Each neural network predictive model \mathbf{E}_b of the ensemble members outputs a Gaussian distribution conditioned on the model input:

$$\mathbf{E}_b(\mathbf{x}_t, \mathbf{u}_t; \theta_b) = \mathcal{N}(\boldsymbol{\mu}_{\theta_b}(\mathbf{x}_t, \mathbf{u}_t), \boldsymbol{\Sigma}_{\theta_b}(\mathbf{x}_t, \mathbf{u}_t)), \quad (1)$$

where $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are the mean and the diagonal covariance parameterized by the model parameter θ_b . Given a collected dataset of transitions $\mathcal{D} \triangleq \{(\mathbf{x}_t, \mathbf{u}_t, \dot{\mathbf{x}}_{t+1})\}_{t=0}^{N-1}$ with the size of N , the probabilistic model can be trained with the Gaussian Negative Log-Likelihood (NLL) loss function:

$$\mathcal{L}_{\text{train}}(\theta_b) = \sum_{t=0}^{N-1} [\boldsymbol{\mu}_{\theta_b} - \dot{\mathbf{x}}_{t+1}]^\top \boldsymbol{\Sigma}_{\theta_b}^{-1} [\boldsymbol{\mu}_{\theta_b} - \dot{\mathbf{x}}_{t+1}] + \log[\det \boldsymbol{\Sigma}_{\theta_b}]. \quad (2)$$

Note that random initialization is applied individually to each ensemble member, resulting in mutually independent initial weights. Consequently, the only difference between the ensemble members is their initial model parameters.

Consequently, the stochastic dynamics function predicting the next state can be interpreted as a Gaussian Mixture Model (GMM) of the B ensemble members:

$$\mathbf{F}(\theta_{1:B}) = \sum_{b=1}^B \pi_b \mathbf{E}_b(\mathbf{x}_t, \mathbf{u}_t; \theta_b), \quad 0 \leq \pi_b \leq 1. \quad (3)$$

Here, we simply use equal weights $\forall \pi_b = \frac{1}{B}$. Then, the next state of the system during iterative trajectory rollouts can be defined as:

$$\mathbf{x}_{t+1} = \mathbf{x}_t + \mathbb{E}[\mathbf{E}_b(\mathbf{x}_t, \mathbf{u}_t; \theta_b)] \Delta t, \quad b \sim \{1, \dots, B\}. \quad (4)$$

B. Model Predictive Active Exploration

To boost the model learning process, we introduce an active exploration strategy in the *training phase*. We encourage the robot to seek automatically to visit currently unforeseen state-action pairs for training data. To this end, ensemble model disagreement D can be used as exploration bonuses [33]. One

possible solution is to measure the empirical variance over the predictive samples $\tilde{\mathbf{x}}_{t+1}$ from the posterior distribution of each ensemble member:

$$\mathfrak{U}(\mathbf{x}_t, \mathbf{u}_t) = \text{Var} \left(\left\{ \tilde{\mathbf{x}}_{t+1} \sim \mathbf{E}_b(\mathbf{x}_t, \mathbf{u}_t) \right\}_{b=1}^B \right), \quad (5)$$

which is referred as *uncertainty sampling*. However, this method cannot distinguish epistemic uncertainty from aleatoric uncertainty. Importantly, the robot should only be inquisitive about the epistemic uncertainty and not the aleatoric uncertainty; otherwise, it will be obscure if this exploration bonus resulted from a lack of knowledge or from inherent and irreducible system stochasticity, such as model ambiguity or sensor noise [32].

Kullback–Leibler (KL) divergence is the most well-known metric for measuring the disagreement between distributions to obtain epistemic uncertainty. However, KL divergence is asymmetric and can only be used to determine the distance between two distributions. Therefore, we employ Jensen–Rényi Divergence (JRD)

$$\text{JRD}(\mathbf{E}_{1:B}; \alpha) \triangleq H_\alpha \left(\sum_{b=1}^B \pi_b \mathbf{E}_b \right) - \sum_{b=1}^B \pi_b H_\alpha(\mathbf{E}_b), \quad (6)$$

where the Rényi entropy [40] of a random variable X and its corresponding density function $p(x)$ is defined as:

$$H_\alpha(X | x \in X) = \frac{1}{1-\alpha} \log \int p(x)^\alpha dx. \quad (7)$$

While an approximation of the JRD can be obtained via Monte-Carlo sampling, it is computationally expensive, thereby precluding real-time implementation. To circumvent this problem, Wang et al. [41] introduce a closed-form JRD with quadratic entropy ($\alpha = 2$) for analytical measurement of the divergence of a GMM:

$$\begin{aligned} D(\mathbf{x}_t, \mathbf{u}_t; \mathbf{E}_{1:B}) = & -\log \left[\frac{1}{B^2} \sum_{i,j} \mathfrak{D}(\mathbf{E}_i, \mathbf{E}_j) \right] + \\ & \frac{1}{B} \sum_i \log [\mathfrak{D}(\mathbf{E}_i, \mathbf{E}_i)], \end{aligned} \quad (8)$$

where

$$\mathfrak{D}(\mathbf{E}_i, \mathbf{E}_j) = \frac{1}{|\Phi|^{\frac{1}{2}}} \exp \left(-\frac{1}{2} \Delta^\top \Phi^{-1} \Delta \right), \quad (9)$$

$\Phi = \Sigma_{\theta_i} + \Sigma_{\theta_j}$ and $\Delta = \mu_{\theta_i} - \mu_{\theta_j}$.

Using this metric, Shyam et al. [33] presented a novel active exploration strategy by measuring the ensemble model disagreement along the future predictions. The proposed method trained a model-free policy to harness this time-varying and intricate exploration bonus. However, due to the sample inefficiency and the inability to adjust the trained policy to achieve modified task-specific objectives, this method has difficulty being used to robotic applications. To make these limitations tractable, we here employ Model Predictive Path Integral (MPPI) control [8] as a policy.

MPPI is the state-of-the-art sampling-based MPC that enables real-time implementation with the aid of Graphic Processing Units (GPUs) due to its parallelizable structure [18]. Denoting by $U = \{\mathbf{u}_0, \dots, \mathbf{u}_{T-1}\}$ with a fixed time horizon T , the MPPI algorithm seeks an optimal control sequence U^* such that:

$$U^* = \underset{U}{\text{argmin}} \mathbb{E} \left[\phi(\mathbf{x}_T) + \sum_{t=0}^{T-1} \mathcal{L}(\mathbf{x}_t, \mathbf{u}_t) \right], \quad (10)$$

where $\phi(\cdot)$ is a state-dependent terminal cost. In general, the running cost with the quadratic control cost with control weight matrix \mathbf{R} takes the form:

$$\mathcal{L}(\mathbf{x}_t, \mathbf{u}_t) = q(\mathbf{x}_t) + \frac{1}{2} \mathbf{u}_t^\top \mathbf{R} \mathbf{u}_t. \quad (11)$$

Note that the state-dependent running cost $q(\cdot)$ can be arbitrary functions, allowed by the sampling-based optimization scheme that is capable of solving non-convex problems. Consequently, we can include the information gain obtained from the model disagreement (8) into the controller’s objective (11) in order to jointly achieve task-dependent cost optimization and active exploration:

$$\mathcal{L}_{\text{active}}(\mathbf{x}_t, \mathbf{u}_t) = q(\mathbf{x}_t) + \frac{1}{2} \mathbf{u}_t^\top \mathbf{R} \mathbf{u}_t - w_D D(\mathbf{x}_t, \mathbf{u}_t), \quad (12)$$

where $w_D > 0$ is a weighting constant. As a result, the robot is encouraged to choose behaviors that lead to exploring uncertain state-action spaces of the dynamic model, while still attempting to minimize the task-specific error.

The above characteristic facilitates the robotics applications of this method in three ways. First, the controller actively seeks states and actions having high epistemic uncertainty, i.e., for which there are little training data, in the vicinity of the set task objective. Thus, the trained model is completely tailored for instantaneous deployment. Second, it is conceivable to determine when the *exploration phase* may be terminated; specifically, it is when the robot optimizes mostly for the task cost. As training data accumulate, the ensemble models will gradually converge toward increasingly similar predictions, resulting in a progressive decrease in model disagreement. Finally, the robot will be able to adapt to the changes in task costs and to solve the modified tasks without re-training.

C. Uncertainty-Aware Deployment

After sufficient training time, the robot may be deployed to perform some designated tasks. In the *deployment phase*, we generally regard that there are no human supervisors and no further dynamics update. Recent research has shown that parallel ensemble neural networks are capable of learning dynamics online [42]. However it is orthogonal to the goal of this paper, and our method also can benefit from this consideration. For successful operations, the robot must avoid unpredictable situations, which can be accomplished by incorporating uncertainty into safety violations, regardless of the source of uncertainty. For instance, a data point may have no disagreement across ensembles due to sufficient training

data (low epistemic uncertainty), but we must still avoid this point if the variance of the predicted posterior distribution is high (high aleatoric uncertainty).

In this phase, we use a variant of the MPPI algorithm called Smooth MPPI (SMPPI) [43]. SMPPI shares the same information-theoretic roots as MPPI, and it also benefits from the structure of parallel trajectory evaluation. However, SMPPI lifts the control variables as derivative actions, so that the noisy sampling is performed on a higher order domain $\dot{U} = \{\dot{\mathbf{u}}_0, \dots, \dot{\mathbf{u}}_{T-1}\}$. Let us denote the resulting optimal control trajectory after the optimization process as \dot{U}_{i+1}^* at MPC iteration step i . Then, the optimal action trajectory U_{i+1}^* , which will be used as the actual commands sent to the robot, is obtained by the integral with respect to the MPC iteration horizon: $U_{i+1}^* = U_i^* + \dot{U}_{i+1}^* \Delta i$, where Δi represents the control updating period of MPC and is commonly equal to Δt . This new action trajectory update law allows the controller to rapidly respond to changing environments while using significantly lower sampling variance, thus alleviating the chattering in resulting commands. Furthermore, we can apply an extra action smoothing cost $\Omega(U)$ along trajectory rollouts:

$$\Omega(U) = \sum_{t=1}^{T-1} (\mathbf{u}_t - \mathbf{u}_{t-1})^\top \boldsymbol{\omega} (\mathbf{u}_t - \mathbf{u}_{t-1}), \quad (13)$$

where $\boldsymbol{\omega} \in \mathbb{R}^{m \times m}$ is the weighting parameter. Such action cost was not eligible in the MPPI baseline with non-affine dynamics because it violates the information theoretic derivation [43]. As a result of these smoothing effects, SMPPI is not suitable for rapid exploration in underactuated continuous control systems but is beneficial for deployment.

To achieve risk-averse planning, taking into account both aleatoric and epistemic uncertainty, we employ the results of *uncertainty sampling* (5) as uncertainty measurements. A simple addition of the quantified uncertainty to the cost function would make the controller uncertainty-aware, as the following form:

$$\mathcal{L}_{\text{naive}}(\mathbf{x}_t, \dot{\mathbf{u}}_t, \mathbf{u}_t) = q(\mathbf{x}_t) + \frac{1}{2} \dot{\mathbf{u}}_t^\top \mathbf{R} \dot{\mathbf{u}}_t + w_1 \mathcal{U}(\mathbf{x}_t, \mathbf{u}_t), \quad (14)$$

where $w_1 > 0$ is a weighting constant. It is important to point out that the measured uncertainty through sampling is not differentiable with respect to state and action. Sampling-based MPC is capable of optimizing any arbitrarily crafted cost functions, making such a form of uncertainty tractable. Similarly, we take use of this trait by imposing an impulse-like penalty where uncertainty exceeds a threshold ξ :

$$\mathcal{L}_{\text{hybrid}}(\mathbf{x}_t, \dot{\mathbf{u}}_t, \mathbf{u}_t) = q(\mathbf{x}_t) + \frac{1}{2} \dot{\mathbf{u}}_t^\top \mathbf{R} \dot{\mathbf{u}}_t + w_1 \mathcal{U}(\mathbf{x}_t, \mathbf{u}_t) + w_2 I\left(\left\{\left|\mathcal{U}(\mathbf{x}_t, \mathbf{u}_t)\right| > \xi\right\}\right), \quad (15)$$

where $w_2 \gg w_1$ is a weighting constant and I is an indicator function. It is an intuitive consideration because we never want to allow the robot to take actions that are completely unpredictable.

Following is the final objective function of SMPPI:

$$\dot{U}^* = \underset{\dot{U}}{\operatorname{argmin}} \mathbb{E} \left[\phi(\mathbf{x}_T) + \Omega(U + \dot{U} \Delta i) + \sum_{t=0}^{T-1} \mathcal{L}_{\text{hybrid}}(\mathbf{x}_t, \dot{\mathbf{u}}_t, \mathbf{u}_t) \right]. \quad (16)$$

D. Implementing Neural Network Vehicle Dynamics

Given that we have made no assumption on the specific robotic platform in our framework, we argue that it is applicable to any kind of robotic system. In this work, we focus on autonomous vehicle and wheeled robot applications to validate our idea. We choose the robot's state and action input based on the dynamic bicycle model since it provides a good trade-off between model accuracy and simplicity for real-time implementation [13]. The dynamic state variable $\mathbf{x} = (v_x, v_y, r)^\top$ consists of longitudinal velocity, lateral velocity, and yaw rate. These states can be easily measured by using GPS and IMU sensors. The action input $\mathbf{u} = (\delta, v_{des})^\top$ consists of the steering angle and desired speed. The desired speed is sent to the low-level Proportional-Integral (PI) controller, which determines the throttle and brake.

Since these state and action variables are simplified representations of the full robot dynamics, there exists irreducible model ambiguity. For example, the roll and pitch motions are not considered in the bicycle model. Automatic gear shifting encompasses non-smooth and time-varying dynamics characteristics, hence increasing the lower bound of model bias. We alleviate this problem by providing the history of state-action pairs to the neural network input so that it can extract contextual information [44, 45]. The history length H must be carefully determined, because as H increases, the state-action space that has to be discovered becomes exponentially larger, making the exploration problem more difficult. We select $H = 4$ in this work through an ablation study in the history length (see Section VI-A). Although LSTM [46] and GRU [47] show slightly better prediction performances than Multi-Layer Perceptron (MLP) (see Section VI-A), we choose MLP since its operations are fully parallelizable.

We take the parallel ensemble MLP implementation by *philipjball* [48] and adapt it for our task. Let us denote the weight and the bias of a layer of an ensemble model \mathbf{E}_b as \mathbf{W}_b and \mathbf{b}_b , respectively. Then, the parameters of all ensemble models can be represented in the form of batches of matrices: $\mathbf{W} = [\mathbf{W}_1^\top, \dots, \mathbf{W}_B^\top]^\top$, $\mathbf{b} = [\mathbf{b}_1, \dots, \mathbf{b}_B]^\top$. First, we initialize the parameters of each ensemble model according to the size of the layer input n_{in} :

$$\mathbf{W}_b, \mathbf{b}_b \sim \mathcal{U} \left(-\sqrt{\frac{1}{n_{\text{in}}}}, +\sqrt{\frac{1}{n_{\text{in}}}} \right), \forall b \in \{1, \dots, B\}, \quad (17)$$

where \mathcal{U} is uniform distribution. While training the model with collected data, the forward propagation of each ensemble model is computed through

$$\mathbf{o}_b = \mathbf{W}_b^\top \mathbf{z} + \mathbf{b}_b, \quad \forall b \in \{1, \dots, B\}, \quad (18)$$

where \mathbf{z} and \mathbf{o} are the input and the output of the layer, respectively. The gradients of the Gaussian NLL loss function (2) can be obtained through automatic differentiation packages such as Pytorch `autograd` [49]. This process is identical to the common way we use MLP. While using the model as the dynamics of the robot, on the other hand, the forward propagation of the entire ensemble model is defined as batch matrix-matrix multiplication and addition:

$$[\mathbf{o}_1, \dots, \mathbf{o}_B]^\top = \mathbf{W}^\top \odot [\mathbf{z}, \dots, \mathbf{z}]^\top \oplus \mathbf{b}, \quad (19)$$

where \odot and \oplus are element-wise operations. This parallelized forward propagation (19) can be implemented using Pytorch `baddbmm` function. Therefore, with the aid of sufficient GPU resources, we can simultaneously predict the next state and obtain uncertainty using PENN dynamics in a fixed-time duration regardless of the number of ensembles. This strategy maximizes the benefits of the parallel sample evaluation property of the MPPI framework.

IV. EXPERIMENTS

Our experimental evaluations address the following key questions: **(Q1)** Can our method explore the state-action space with sufficient sample efficiency (see Section IV-A)? **(Q2)** Can the PENN taught with active exploration be used for long horizon planning, even if the task objective has been modified after training (see Section IV-B)? **(Q3)** Can our method be turned directly into an uncertainty-aware controller with minimal modification (see Section IV-C)? **(Q4)** Is our method scalable to other robot platforms (see Section IV-D)?

We evaluate **Q1-Q3** in a high-fidelity vehicle dynamics simulator - IPG CarMaker. We build the modeling and simulation system by integrating the algorithms using ROS 2 [50]. A Land Rover Defender 110 is used as the control vehicle. The low-level PI controller runs at 100 Hz, and the control parameters K_P, K_I are set to $\{0.35, 1.0\}$ for throttle and $\{0.18, 0.3\}$ for brake. To simulate sensor noise, we add 10% of i.i.d. Gaussian noise to each dynamic state observation in \mathbf{x} . MPPI is used in the active exploration experiments **(Q1)**, and SMPPI is used in the deployment experiments **(Q2-Q3)**. Table I lists the parameters of two controllers, while the remaining parameters not specified in this paper are taken from the SMPPI implementation [43].

TABLE I: Control parameters of MPPI and SMPPI. Note that the action smoothing cost with ω only applies to SMPPI.

Parameters	MPPI [18]	SMPPI [43]
Δt	0.1 s	0.1 s
Δi	0.1 s	0.1 s
K	5,000	5,000
T	10	35
Sampling Variance	Diag(4.0, 3.0)	Diag(1.6, 0.4)
ω	N/A	Diag(0.4, 0.1)

We use a four-layer MLP with the hidden layer sizes $\{40, 80, 120, 40\}$ throughout all experiments. Rectified Linear

Unit (ReLU) [51] is used as activation functions. Similar to [27], we use five ensembles ($B = 5$) for real-time implementation.

A. Model Predictive Active Exploration

Experimental Setup. In this experiment, we evaluate the effectiveness of the active exploration strategy **(Q1)**. We let the robot seek for training data on a large flat ground with 0.7 friction coefficient. The task objective is to maintain the 50 km/h target speed (v_{target}) while staying inside the space boundary. Accordingly, the task-dependent running cost $q(\mathbf{x}_t)$ is defined as follows:

$$\begin{aligned} q(\mathbf{x}_t) &= \alpha_1 \text{Track}(\mathbf{x}_t) + \alpha_2 \text{Speed}(\mathbf{x}_t), \\ \text{Track}(\mathbf{x}_t) &= (0.9)^t 10000 \mathbf{M}(p_x, p_y), \\ \text{Speed}(\mathbf{x}_t) &= \left(\sqrt{v_x^2 + v_y^2} - v_{\text{target}} \right)^2. \end{aligned} \quad (20)$$

$\alpha_{(\cdot)}$ are empirically tuned weighting constants. The track cost imposes a hard penalty to prevent collisions. The given map \mathbf{M} indicates whether the robot position (p_x, p_y) is at the outside of the space boundary.

We assume that there are human monitors who supervise the *exploration phase*. They disengage the robot's autonomy before the robot reaches the space boundary or before the robot enters dangerous situations. We empirically assume this period as 30 s. While the human supervisors setting up the robot for a new trial, the PENN is trained with the collected data offline. We refer to this process as one *iteration* for exploration.

Methods Evaluated. We evaluate the following methods including ours:

- 1) *Ours*: a model predictive active exploration strategy, which optimizes the objective function of (12). The epistemic uncertainty measured by JRD (8) serves as the information gain. We set the information gain weighting constant w_D as 100.
- 2) *Uncertainty Sampling (US)*: an uncertainty sampling counter-part of our method, which uses the uncertainty quantified through sampling over each ensemble posterior distribution as the information gain. The substitution of JRD (8) with uncertainty sampling (5) in (12) is used as its running cost. We use the same weighting $w_D = 100$ as our method since the quantified information gain in both methods is of similar magnitude in our implementation.

3) *Jensen-Rényi Divergence Reactive Exploration (JDRX)*: a reactive counter-part of our method [33]. We simply implement this based on our method by setting the planning horizon to $T = 1$, which explores greedily without planning based on the experience collected so far.

4) *Random Noise (RN)*: a random exploration strategy, which injects noise in the action input. Its objective function is (11) and there is no information gain. We sample random noise from the Gaussian distribution with the same variances as the sampling distribution used by MPPI.

All the compared methods are evaluated with the same control parameters and the same model training setting with a fixed random seed. While training the model offline, we

randomly split the collected data \mathcal{D} into the training and test sets with a 7:3 ratio. The PENN is trained using the Adam optimizer with a learning rate of 0.001. We save the model that shows the best test sets performance during 10 epochs, and use this model as the robot dynamics \mathbf{F} in the next iteration.

Experimental Results.

We first evaluate how much of the state-action space each method has covered. We compute the convex hull [52] with the collected state and action data $\{\mathbf{x}_t, \mathbf{u}_t\}_{t=0}^{N-1}$ at each iteration, and analyze the volume enlargement of the convex hulls [53] (see Fig. 2). The results indicate that JDRX cannot explore the state-action space efficiently. Also, the volumes of ours and RN show comparable magnitudes while consistently exceeding those of US.

We further analyze the collected data points in detail. Similar to as Spielberg et al. [54], we focus on the state envelope of sideslip β and yaw rate r , since precise motion predictions with large sideslip angles and rotational speed are crucial to the success of high-speed driving on sharp curves. The results are shown in Fig 3. JDRX is not able to collect driving data with high yaw rates. This is because it reacts immediately to the trivial information gain near the current state and is unable to plan for higher information gain in the future. On the other hand, ours always covers the largest state space among the compared methods. This is because ours can plan over a longer horizon to gather drifting data by ignoring the aleatoric uncertainty exudes from plain driving. Our method has collected a substantial amount of data with β around 0.2 rad (approximately 11.5 °), and even data with β around 0.3 rad (approximately 17.2 °). Lastly, we stress that the data collected by RN is not helpful for achieving the driving tasks, despite it showing a similar enlargement in the convex hull to our method. We will further discuss this in Section IV-B.

B. Direct Deployment

Experimental Setup. We take the resulted PENNs of active exploration as the robot dynamics and directly evaluate their performances in long horizon planning (Q2). We load the best models at each iteration and monitor the control performance improvements of the PENNs of the compared methods. The performances are evaluated in a race track with three moderate curves and five sharp curves (see Fig. 4a). We set the friction coefficient of the track to 0.7.

In this experiment, the objective of the robot is to drive through the race track at a speed as close to v_{target} as possible, while staying inside of the track. We modify the task-

dependent cost $q(\mathbf{x}_t)$ as follows:

$$q(\mathbf{x}_t) = \alpha_1 \text{Track}(\mathbf{x}_t) + \alpha_2 \text{Speed}(\mathbf{x}_t) + \alpha_3 \text{Stable}(\mathbf{x}_t). \quad (21)$$

The track cost and the speed cost are nearly identical to those in the exploration phase. The given map \mathbf{M} is modified accordingly to the new race track. The target speed v_{target} is adjusted to 45 km/h. We add the stabilizing cost that imposes a hard penalty on large sideslip angles β to prevent the robot from losing its stability:

$$\begin{aligned} \text{Stable}(\mathbf{x}_t) &= 10000 I(\{|\beta| > 0.3\}), \\ \beta &= -\arctan\left(\frac{v_y}{\|v_x\|}\right). \end{aligned} \quad (22)$$

We set the planning horizon to $T = 35$, which is 3.5 s, to compare the performances in long horizon planning. Note that no information gain is provided to the objective function, so the controllers are uncertainty-neutral.

Experimental Results. The results are shown in Fig. 4. At the sharpest corners of #3, #7, and #8, in order to drive through them without losing significant speed, the vehicle should gently slow down before entering the corners and drift continuously. If the vehicle cannot perform drifting maneuvers, it will nearly stop in the middle of the corners or even collide with the track boundary. Moreover, if the predictions of the dynamics are not accurate, the vehicle may enter a saturated drifting region and become uncontrollable.

First, we count the number of times that each method completes the whole track without driving outside of the given boundary (see Fig. 4b). JDRX is not able to complete the whole lap. Reactive exploration limits the robot to acquiring only insignificant information gain in the vicinity of the current state. As a result, it fails to collect high-speed driving data for training and to produce accurate predictions in states with high speeds. RN shows improvement as it collects more data through repeated iterations. However, it only relies on coincidental actions that lead to collecting novel data. Our method achieves the best results among the compared methods, and its performance steadily improves as more data is collected.

Interestingly, the performance of US improves in the early stages, but after 100 iterations, its performance continuously decreases. It implies that collecting more data does not necessarily improve the prediction performance. Since US cannot separate the source of uncertainty, it may repeatedly collect the same data showing naive motions. Therefore, the model is trained to make increasingly better predictions on these superfluous data while neglecting the data representing important driving characteristics.

We also analyze the control performance in terms of speed (see Fig. 4c) and stability (see Fig. 4d) during each trial. Our method has effectively regulated the speed cost compared to the other methods. The results also show that RN violates the stabilizing constraints far more frequently than our method. It demonstrates that the collected data from RN do not effectively contribute to achieving the given task, since they omit crucial training data such as high-speed drifting




Fig. 2: The volumes of the convex hulls along exploration iterations.




Fig. 3: The scatter plots of collected data during active exploration. The data are from 10, 30, 100, and 300 iterations, respectively. During all iterations, ours using JRD information gain covers the largest state spaces of sideslip angle and yaw rate compared to other methods.

maneuvers. Only our method successfully drives through the race track at high speeds without losing stability.

C. Uncertainty-Aware Deployment

Experimental Setup. We assumed that the *exploration phase* was performed in a large open space under human supervision. However, in *deployment phase*, there is no such monitoring to guarantee safety. In this experiment, we validate that our active exploration policy can be transferred to an uncertainty-aware controller with minimal modification.

To highlight the influence of the uncertainty-aware control strategy, we augment the dynamics to be terrain-aware by feeding the elevation map to the neural network. We assume that a 2.5D elevation map with a 0.1 m grid size is provided by one of the existing methods that can generate these maps from raw sensor input in real time [55, 56]. The map encoder network consists of an 8-channel convolutional layer and a 4-channel convolutional layer, both with a kernel size of 3 and a stride of 2 [57]. The output of the encoder is concatenated with the history of state-action pairs and fed into the original MLP model that had been used throughout the previous experiments. We use the data collected by our method, which shows the best performance in direct deployment experiments, to train the augmented dynamics. During training, the elevation map is randomly generated from the Gaussian distribution of $\mathcal{N}(0, 0.2)$.

Following our prior work [58], we design a realistic off-road environment with various components including large bumps and randomly patterned rough terrains (see Fig. 5). We also add non-traversable elements such as trees, bushes, and wooden pillars, and they are included in the given map \mathbf{M} for collision avoidance.

Methods Evaluated. We take the same task-dependent cost function (21) used in the direct deployment experiments. The target speed v_{target} is adjusted to 30 km/h. The key difference between these two deployment tasks is that we here include uncertainty in the control objective. We evaluate the following uncertainty-averse strategies:

1) *Uncertainty-Neutral (UN)*: a uncertainty neutral strategy that does not take uncertainty into account while planning. It

uses the same objective function as the one used in direct deployment experiments.

2) *Naive Penalty (NP)*: a simple uncertainty-aware strategy that penalizes a linear cost on quantified uncertainty. It uses (14) as the objective function with $w_1 = 1000$.

3) *Hybrid*: an uncertainty-aware strategy that has both soft and hard costs. It uses (15) as the objective function with $w_1 = 1000$, $w_2 = 10000$, and $\xi = 0.15$. Note that the hard cost penalizes states that are estimated to have uncertainty greater than ξ with the same magnitude as a collision with the track boundary.

4) *Conservative*: a conservative strategy that controls the vehicle as stable as possible. It uses (15) as the objective function with $w_1 = 0$, $w_2 = 10000$, and $\xi = 0.02$.

Experimental Results. We visualize the trajectories taken by the compared methods in Fig. 5. To illustrate the vehicle’s stability, we display the rotational impacts exerted on the vehicle during driving onto the trajectories. We simply use angular momentum, obtained by the squared sum of roll rate and pitch rate. We normalize the angular momentum by the maximum angular momentum value across all trials. Therefore, a trajectory with a darker color represents that the vehicle is more unstable. The angular and vertical motions of the vehicle during experiments are shown in Table II for quantitative evaluation.

Although UN can control the vehicle at the reference speed while avoiding collisions with obstacles and staying on track, it fails to circumvent uncertain regions because it is blind to uncertainty. It drives over the bumps at Scenario #1 and Scenario #5, eventually resulting in overturning. NP has opted for terrain with less uncertainty and, as a result, has chosen moderate terrain rather than the high bumps (see Scenario #1 and Scenario #5). However, it fails to circumvent the largest bump at Scenario #3. On the other hand, Hybrid decides to slow down before the bump and take a detour to satisfy the safety criteria, despite the speed penalty. Conservative is a reasonable option in practical applications because it minimizes the damage to the robot. It takes the most secure trajectory as can be seen at Scenario #4 and Scenario #5 due to the low uncertainty threshold. The trajectory taken by

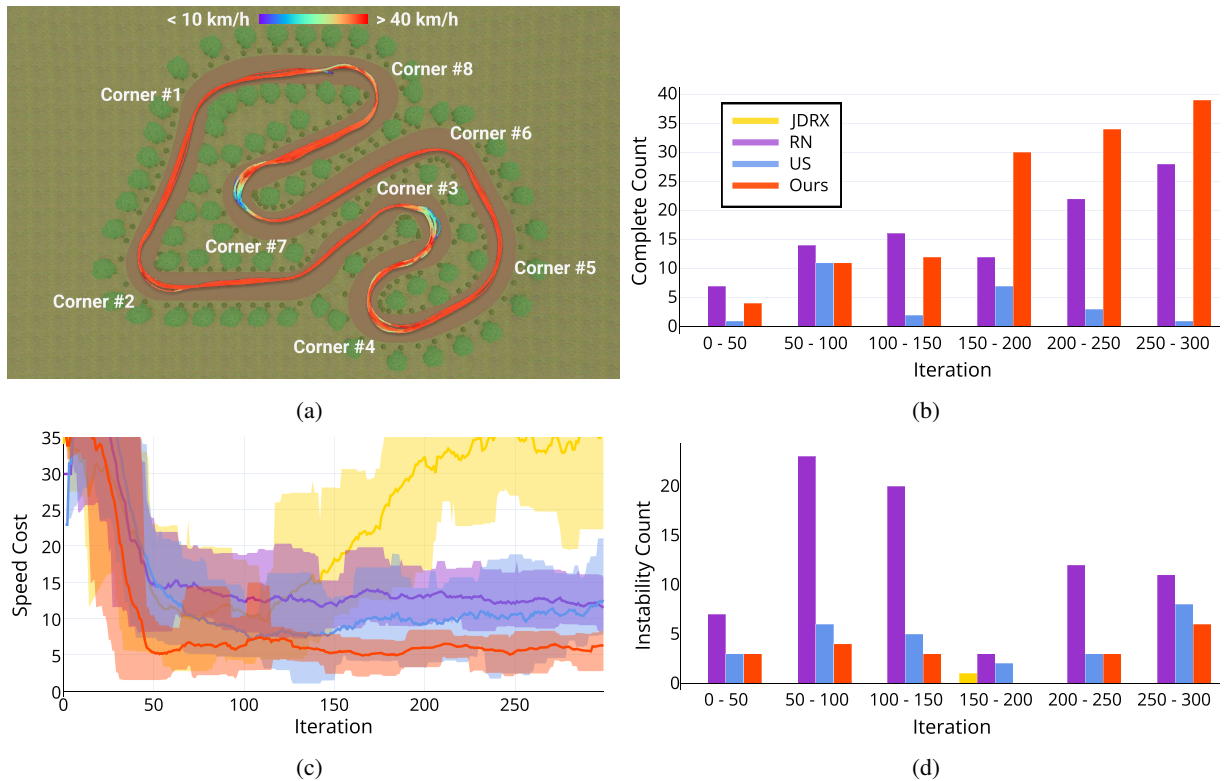


Fig. 4: (a) The race track designed in the IPG CarMaker simulator for direct deployment experiments. We visualize the vehicle trajectory taken by our method at 300 iterations for driving 10 laps in a counter-clockwise direction. (b) The number of times that each method completes the whole lap during every 50 iterations. (c) The average speed cost for each saved model along exploration iterations. The shaded areas denote a 95% confidence interval. (d) The number of trials of each method that violates the stabilizing constraints at least once, i.e., having larger sideslip angles than 0.3 rad, during every 50 iterations.

TABLE II: The average and maximum motions of the vehicle across all trials. The vertical motions include vertical velocity and vertical acceleration. The angular motions include velocities and accelerations of roll and pitch, respectively. Hybrid and Conservative experience lower terrain impacts than UN and NP by a large margin.

#	Method	Vertical Vel. [m/s]		Vertical Acc. [m/s ²]		Roll Rate [rad/s]		Pitch Rate [rad/s]		Roll Acc. [rad/s ²]		Pitch Acc. [rad/s ²]	
		Mean	Max	Mean	Max	Mean	Max	Mean	Max	Mean	Max	Mean	Max
1	Uncertainty-Neutral (UN)	2.07	28.23	0.13	1.09	0.28	2.46	0.43	1.10	2.22	30.67	1.64	12.78
2	Naive Penalty (NP)	1.97	23.30	0.05	0.38	0.33	2.06	0.44	2.22	2.73	33.97	1.80	8.02
3	Hybrid	1.42	9.37	0.04	0.20	0.21	1.56	0.39	1.15	1.85	18.37	1.28	5.91
4	Conservative	1.40	10.60	0.04	0.22	0.14	1.60	0.42	1.11	1.16	22.23	1.46	6.93

Conservative also shows that the vehicle rarely experienced terrain impacts except at the inevitable bump at Scenario #2. Hybrid shows a good balance between safety and agility. It avoids regions with high degrees of uncertainty, which might irreversibly damage or destroy the vehicle, while navigating the terrains with low levels of uncertainty at the desired speed. Table II demonstrates that the methods with hard uncertainty penalty show a significantly lower degree of undesirable vertical motions and angular motions compared to the other methods.

D. Experiments on Wheeled Robots

In this section, we conduct a sanity check to ensure that our framework is scalable to other types of robot platforms.

We use a 1:5 scale wheeled robotics testbed as the novel target system. We examine our framework in two robotics simulators: Gazebo [59] and Nvidia Isaac Sim [60]. In each simulator, we conduct the active exploration experiments (Q1) and direct deployment experiments (Q2) with the best performing exploration strategy. Note that we do not modify any settings in our framework except for adjusting the target speed according to the robot’s specifications ($v_{\text{target}} = 15 \text{ km/h}$) in both exploration and deployment phases. Other experimental settings are identical to those described in previous sections, including parameters of neural networks, controllers, and objective functions. Fig. 6 shows the simulated environments and the resulting trajectories in deployment tests taken by

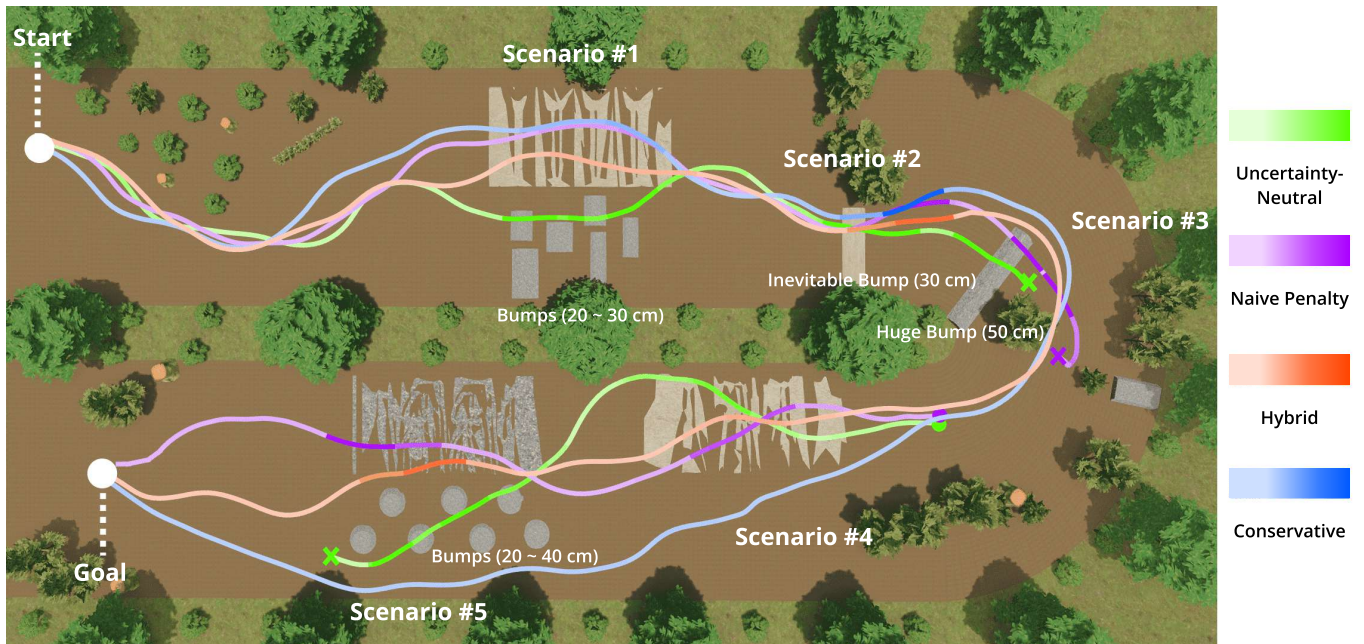


Fig. 5: Visualization of uncertainty-aware navigation results on the vehicle simulator. We display the rotational impacts exerted on the vehicle onto the trajectories. The maximum value of rotational impact during a three-second window is used for visual clarity.

the robots. The results suggest that our framework has the potential to be applied to other robot platforms with continuous state and action space, actively exploring training data, and safely controlling the robot.

V. CONCLUSION

In this paper, we presented a unified model-based reinforcement learning framework with dynamics learning that bridges active exploration and uncertainty-aware deployment in the robotics control domain. The dynamics is learned using a fully parallelized probabilistic ensemble neural network that is sensitive to uncertainty. For active exploration, the epistemic uncertainty can be quantified by measuring the ensemble disagreement via Jensen-Rényi Divergence. Two opposing task objectives for exploration and deployment can be optimized by the state-of-the-art sampling-based MPC. Our extensive experiments demonstrate that the unified framework can be applied to both autonomous vehicles and wheeled robots. Our framework shows promising results for both exploration and deployment, in that the robot efficiently collects useful training data samples that are essential to achieving the task, and successfully avoids uncertain regions by imposing an uncertainty penalty. We hope that this work will serve as a stepping stone towards more general learning-based robotic applications under uncertainty.

VI. APPENDIX

A. Additional Details for Learned Dynamics

To perform an ablation study on the neural network dynamics, we collect a human-controlled driving dataset. We build a race track in the IPG CarMaker simulator, modeled after a kart

circuit located in Kirchlengern, Germany [43]. The length of the track is 1016 m and it has two moderate curves and four sharp curves. The overall structure of the track is similar to that in Fig. 6c. The driving data collected on the track consists of 1) zig-zag driving at low speeds, 2) high-speed driving, and 3) sliding maneuvers, in both clockwise and counter-clockwise directions. We collected a total of 70 minutes of data at a rate of 10 Hz. We split the data into 70% for training and 30% for testing after randomizing the data to remove temporal correlations.

We evaluate the test performance of the compared models with respect to each state variable in $\mathbf{x} = (v_x, v_y, r)^\top$, which are the longitudinal velocity v_x , the lateral velocity v_y , and the yaw rate r . We use the Root Mean Square Error (RMSE) as the performance metric. We also display the average test RMSE for the convenience of comparison.

Model Comparison. First, we compare the pure MLP model and the Recurrent Neural Network (RNN) model. In this experiment, we use a deterministic model without ensemble and use L2 loss function. The other parameters of the MLP model follow those mentioned in Section IV. In RNN models, we replace the first layer of the MLP model as GRU [47] and LSTM [46], respectively, with the same hidden layer size as the MLP model. The best test performances of each model during 1000 epochs are shown in Table III. The RNN models show slightly better performances than the MLP model. However, we choose MLP to benefit from the parallelizable structure of the ensemble MLP model.

History Comparison. In this experiment, we use PENN model as described in Section III-D. We evaluate the perfor-




Fig. 6: Additional experiments using a 1:5 scale wheeled robotics testbed in (a)-(c) Gazebo and (d)-(f) Nvidia Isaac Sim. (a),(d) The simulated environments during the exploration phase. (b),(e) The simulated environments during the deployment phase. (c),(f) The vehicle trajectories taken by the robot at 300 iterations for driving 10 laps in a counter-clockwise direction.

TABLE III: A comparison between the MLP model and the RNN models. The prediction performance on the test set is evaluated with the RMSE.

Model	MLP	GRU [47]	LSTM [46]
Total	0.0802	0.0721	0.0701
v_x [m/s]	0.0990	0.0861	0.0843
v_y [m/s]	0.06510	0.0587	0.0576
r [rad/s]	0.0707	0.0660	0.0647

manances of the PENNs with different lengths of state-action history H . The results are shown in Table IV. The lengths of 3, 4, and 5 produce the best results among the compared models. When the history length exceeds 5, the prediction performance generally decreases. In accordance with prior literature employing the same strategy [44, 45, 54, 61], we determine the history length to be 4.

B. Sim-to-Real Transfer

We successfully transferred our algorithm to real-world settings for uncertainty-aware deployment tasks. We integrated our algorithm with global path planning and online traversability map generation using a LiDAR sensor. These experiments were conducted on our off-road testbeds. The experimental results can be found on our project page: <https://taekyung.me/rss2023-bridging>.

VII. ACKNOWLEDGMENT

This work was supported by the Korean Government (2023).

REFERENCES

- [1] J. Morimoto and K. Doya, “Acquisition of stand-up behavior by a real robot using hierarchical reinforcement learning,” *Robotics and Autonomous Systems*, vol. 36, no. 1, pp. 37–51, 2001. [ii](#)
- [2] X. B. Peng, G. Berseth, K. Yin, and M. Van De Panne, “DeepLoco: dynamic locomotion skills using hierarchical deep reinforcement learning,” *ACM Transactions on Graphics*, vol. 36, no. 4, pp. 41:1–41:13, 2017. [ii](#)
- [3] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, “Learning agile and dynamic motor skills for legged robots,” *Science Robotics*, vol. 4, no. 26, p. eaau5872, 2019. [ii](#)
- [4] A. Kumar, Z. Fu, D. Pathak, and J. Malik, “Rma: Rapid motor adaptation for legged robots,” in *Robotics: Science and Systems (RSS)*, 2021. [ii](#)
- [5] P. Abbeel, A. Coates, M. Quigley, and A. Ng, “An application of reinforcement learning to aerobatic helicopter flight,” in *Neural Information Processing Systems (NeurIPS)*, 2007, pp. 1–8. [ii](#)
- [6] M. Deisenroth and C. E. Rasmussen, “PILCO: A model-based and data-efficient approach to policy search,” in *International Conference on Machine Learning (ICML)*, 2011, pp. 465–472. [ii](#)
- [7] P. Englert, A. Paraschos, J. Peters, and M. P. Deisenroth, “Model-based imitation learning by probabilistic trajectory matching,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2013, pp. 1922–1927. [ii](#)
- [8] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou, “Information theoretic MPC for model-based reinforcement learning,”

TABLE IV: A comparison between different lengths of state-action history H . The prediction performance on the test set is evaluated with the RMSE.

History Length	1	2	<u>3</u>	<u>4</u>	<u>5</u>	6	7	8	9	10
Total	0.0623	0.0589	0.0423	0.0453	0.0437	0.0667	0.0852	0.0528	0.0523	0.0511
v_x [m/s]	0.0779	0.0662	0.0548	0.0610	0.0593	0.0826	0.0982	0.0734	0.0742	0.0709
v_y [m/s]	0.0637	0.0691	0.0373	0.0389	0.0373	0.0701	0.1024	0.0429	0.0408	0.0419
r [rad/s]	0.0389	0.0357	0.0319	0.0305	0.0284	0.0399	0.0404	0.0336	0.0321	0.0327

- in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017, pp. 1714–1721. [ii](#), [iv](#)
- [9] G. Kahn, A. Villafior, V. Pong, P. Abbeel, and S. Levine, “Uncertainty-Aware Reinforcement Learning for Collision Avoidance,” in *arXiv preprint arXiv:1702.01182*, 2017. [ii](#), [iii](#)
- [10] A. Nagabandi, K. Konolige, S. Levine, and V. Kumar, “Deep dynamics models for learning dexterous manipulation,” in *Conference on Robot Learning (CoRL)*, 2019, pp. 1101–1112. [ii](#)
- [11] P. Wu, A. Escontrela, D. Hafner, K. Goldberg, and P. Abbeel, “DayDreamer: World Models for Physical Robot Learning,” in *Conference on Robot Learning (CoRL)*, 2022. [ii](#)
- [12] D. Nguyen-Tuong and J. Peters, “Using model knowledge for learning inverse dynamics,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2010, pp. 2677–2682. [ii](#)
- [13] J. Kabzan, L. Hewing, A. Liniger, and M. N. Zeilinger, “Learning-based model predictive control for autonomous racing,” *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3363–3370, 2019. [ii](#), [v](#)
- [14] M. Lutter, C. Ritter, and J. Peters, “Deep Lagrangian Networks: Using Physics as Model Prior for Deep Learning,” in *International Conference on Learning Representations (ICLR)*, 2019. [ii](#)
- [15] M. O’Connell, G. Shi, X. Shi, K. Azizzadenesheli, A. Anandkumar, Y. Yue, and S.-J. Chung, “Neural-Fly enables rapid learning for agile flight in strong winds,” *Science Robotics*, vol. 7, no. 66, p. eabm6597, 2022. [ii](#)
- [16] J. Kober, J. A. Bagnell, and J. Peters, “Reinforcement learning in robotics: A survey,” *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1238–1274, 2013. [ii](#)
- [17] T. Kim, H. Lee, and W. Lee, “Physics Embedded Neural Network Vehicle Model and Applications in Risk-Aware Autonomous Driving Using Latent Features,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 4182–4189. [ii](#)
- [18] 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. [ii](#), [iv](#), [vi](#)
- [19] Y. Yang, K. Caluwaerts, A. Iscen, T. Zhang, J. Tan, and V. Sindhwani, “Data Efficient Reinforcement Learning for Legged Robots,” in *Conference on Robot Learning (CoRL)*, 2019, pp. 1–10. [ii](#)
- [20] C. Williams and C. Rasmussen, “Gaussian Processes for Regression,” in *Neural Information Processing Systems (NeurIPS)*, vol. 8. MIT Press, 1995. [ii](#)
- [21] D. Nguyen-Tuong and J. Peters, “Model learning for robot control: a survey,” *Cognitive Processing*, vol. 12, no. 4, pp. 319–340, 2011. [ii](#)
- [22] S. Levine and P. Abbeel, “Learning Neural Network Policies with Guided Policy Search under Unknown Dynamics,” in *Neural Information Processing Systems (NeurIPS)*, vol. 27, 2014. [ii](#)
- [23] M. Milanese and A. Vicino, “Optimal estimation theory for dynamic systems with set membership uncertainty: An overview,” *Automatica*, vol. 27, no. 6, pp. 997–1009, 1991. [ii](#)
- [24] A. Parsi, A. Iannelli, and R. S. Smith, “Active exploration in adaptive model predictive control,” in *IEEE Conference on Decision and Control (CDC)*, 2020, pp. 6186–6191. [ii](#)
- [25] S. Vijayakumar, A. D’souza, and S. Schaal, “Incremental online learning in high dimensions,” *Neural Computation*, vol. 17, no. 12, pp. 2602–2634, 2005. [ii](#)
- [26] G. R. Williams, B. Goldfain, K. Lee, J. Gibson, J. M. Rehg, and E. A. Theodorou, “Locally Weighted Regression Pseudo-Rehearsal for Adaptive Model Predictive Control,” in *Conference on Robot Learning (CoRL)*, 2019, pp. 969–978. [ii](#)
- [27] K. Chua, R. Calandra, R. McAllister, and S. Levine, “Deep Reinforcement Learning in a Handful of Trials using Probabilistic Dynamics Models,” in *Neural Information Processing Systems (NeurIPS)*, 2018. [iii](#), [vi](#)
- [28] M. Vlastelica, S. Blaes, C. Pinneri, and G. Martius, “Risk-Averse Zero-Order Trajectory Optimization,” in *Conference on Robot Learning (CoRL)*, 2021. [iii](#)
- [29] S. Thrun, “Efficient Exploration In Reinforcement Learning,” Carnegie Mellon University, Pittsburgh, PA, Tech. Rep. CMU-CS-92-102, January 1992. [iii](#)
- [30] R. Martinez-Cantin, M. Lopes, and L. Montesano, “Body schema acquisition through active learning,” in *IEEE International Conference on Robotics and Automation (ICRA)*, 2010, pp. 1860–1866. [iii](#)
- [31] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell, “Curiosity-driven Exploration by Self-supervised Predic-

- tion,” in *International Conference on Machine Learning (ICML)*, 2017, pp. 2778–2787. [iii](#)
- [32] M. Schultheis, B. Belousov, H. Abdulsamad, and J. Peters, “Receding Horizon Curiosity,” in *Conference on Robot Learning (CoRL)*, 2019, pp. 1278–1288. [iii](#), [iv](#)
- [33] P. Shyam, W. Jaśkowski, and F. Gomez, “Model-based active exploration,” in *International Conference on Machine Learning (ICML)*, 2019, pp. 5779–5788. [iii](#), [iv](#), [vi](#)
- [34] R. Sekar, O. Rybkin, K. Daniilidis, P. Abbeel, D. Hafner, and D. Pathak, “Planning to explore via self-supervised world models,” in *International Conference on Machine Learning (ICML)*, 2020, pp. 8583–8592. [iii](#)
- [35] B. Bucher, K. Schmeckpeper, N. Matni, and K. Daniilidis, “An Adversarial Objective for Scalable Exploration,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 2670–2677. [iii](#)
- [36] S. Behtle, Y. Lin, A. Rai, L. Righetti, and F. Meier, “Curious iLQR: Resolving Uncertainty in Model-based RL,” in *Conference on Robot Learning (CoRL)*, 2019, pp. 162–171. [iii](#)
- [37] T. Yu, G. Thomas, L. Yu, S. Ermon, J. Zou, S. Levine, C. Finn, and T. Ma, “MOPO: Model-based Offline Policy Optimization,” in *Neural Information Processing Systems (NeurIPS)*, 2020. [iii](#)
- [38] S. J. Wang, S. Triest, W. Wang, S. Scherer, and A. Johnson, “Rough Terrain Navigation Using Divergence Constrained Model-Based Reinforcement Learning,” in *Conference on Robot Learning (CoRL)*, 2021. [iii](#)
- [39] J. Buckman, D. Hafner, G. Tucker, E. Brevdo, and H. Lee, “Sample-Efficient Reinforcement Learning with Stochastic Ensemble Value Expansion,” in *Neural Information Processing Systems (NeurIPS)*, 2018. [iii](#)
- [40] A. Rényi, “On measures of entropy and information,” in *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability*, 1961, pp. 547–562. [iv](#)
- [41] F. Wang, T. Syeda-Mahmood, B. C. Vemuri, D. Beymer, and A. Rangarajan, “Closed-Form Jensen-Renyi Divergence for Mixture of Gaussians and Applications to Group-Wise Shape Registration,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, vol. 5761, 2009, pp. 648–655. [iv](#)
- [42] T. Z. Jiahao, K. Y. Chee, and M. A. Hsieh, “Online Dynamics Learning for Predictive Control with an Application to Aerial Robots,” in *Conference on Robot Learning (CoRL)*, 2022. [iv](#)
- [43] T. Kim, G. Park, K. Kwak, J. Bae, and W. Lee, “Smooth Model Predictive Path Integral Control Without Smoothing,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10406–10413, 2022. [v](#), [vi](#), [x](#)
- [44] N. A. Spielberg, M. Brown, N. R. Kapania, J. C. Kegelman, and J. C. Gerdes, “Neural network vehicle models for high-performance automated driving,” *Science Robotics*, vol. 4, no. 28, 2019. [v](#), [xi](#)
- [45] T. Kim, H. Lee, S. Hong, and W. Lee, “TOAST: Trajectory Optimization and Simultaneous Tracking Using Shared Neural Network Dynamics,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 9747–9754, 2022. [v](#), [xi](#)
- [46] S. Hochreiter and J. Schmidhuber, “Long Short-Term Memory,” *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, 1997. [v](#), [x](#), [xi](#)
- [47] K. Cho, B. van Merriënboer, D. Bahdanau, and Y. Bengio, “On the Properties of Neural Machine Translation: Encoder-Decoder Approaches,” in *arXiv preprint arXiv:1409.1259*, 2014. [v](#), [x](#), [xi](#)
- [48] “Implement truly parallel ensemble layers by philipjball.” [Online]. Available: <https://github.com/pytorch/pytorch/issues/54147> [v](#)
- [49] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, “Automatic differentiation in PyTorch,” 2017. [vi](#)
- [50] S. Macenski, T. Foote, B. Gerkey, C. Lalancette, and W. Woodall, “Robot Operating System 2: Design, architecture, and uses in the wild,” *Science Robotics*, vol. 7, no. 66, p. eabm6074, 2022. [vi](#)
- [51] V. Nair and G. E. Hinton, “Rectified Linear Units Improve Restricted Boltzmann Machines,” in *International Conference on Machine Learning (ICML)*, 2010. [vi](#)
- [52] C. B. Barber, D. P. Dobkin, and H. Huhdanpaa, “The quickhull algorithm for convex hulls,” *ACM Transactions on Mathematical Software*, vol. 22, no. 4, pp. 469–483, 1996. [vii](#)
- [53] X. Liu, H. Jia, Y. Wen, Y. Hu, Y. Chen, C. Fan, Z. HU, and Y. Yang, “Towards Unifying Behavioral and Response Diversity for Open-ended Learning in Zero-sum Games,” in *Neural Information Processing Systems (NeurIPS)*, vol. 34, 2021, pp. 941–952. [vii](#)
- [54] N. A. Spielberg, M. Brown, and J. C. Gerdes, “Neural Network Model Predictive Motion Control Applied to Automated Driving With Unknown Friction,” *IEEE Transactions on Control Systems Technology*, vol. 30, no. 5, pp. 1934–1945, 2022. [vii](#), [xi](#)
- [55] P. Fankhauser, M. Bloesch, and M. Hutter, “Probabilistic Terrain Mapping for Mobile Robots With Uncertain Localization,” *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3019–3026, 2018. [viii](#)
- [56] T. Miki, L. Wellhausen, R. Grandia, F. Jenelten, T. Homberger, and M. Hutter, “Elevation Mapping for Locomotion and Navigation using GPU,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 2273–2280. [viii](#)
- [57] H. Karnan, K. S. Sikand, P. Atreya, S. Rabiee, X. Xiao, G. Warnell, P. Stone, and J. Biswas, “VI-IKD: High-Speed Accurate Off-Road Navigation using Learned Visual-Inertial Inverse Kinodynamics,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 3294–3301. [viii](#)

- [58] J. Seo, T. Kim, K. Kwak, J. Min, and I. Shim, “ScaTE: A Scalable Framework for Self- Supervised Traversability Estimation in Unstructured Environments,” *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 888–895, 2023. [viii](#)
- [59] N. Koenig and A. Howard, “Design and use paradigms for Gazebo, an open-source multi-robot simulator,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, vol. 3, 2004, pp. 2149–2154. [ix](#)
- [60] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, and G. State, “Isaac Gym: High Performance GPU-Based Physics Simulation For Robot Learning,” in *arXiv preprint arXiv:2108.10470*, 2021. [ix](#)
- [61] N. O. Lambert, D. S. Drew, J. Yaconelli, S. Levine, R. Calandra, and K. S. J. Pister, “Low-Level Control of a Quadrotor With Deep Model-Based Reinforcement Learning,” *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 4224–4230, 2019. [xi](#)