

# Bridging the Sim-to-Real Gap with Bayesian Inference

Jonas Rothfuss<sup>\*,1</sup>, Bhavya Sukhija<sup>\*,1</sup>, Lenart Treven<sup>\*,1</sup>,  
Florian Dörfler<sup>2</sup>, Stelian Coros<sup>1</sup>, Andreas Krause<sup>1</sup>

**Abstract**—We present SIM-FSVGD for learning robot dynamics from data. As opposed to traditional methods, SIM-FSVGD leverages low-fidelity physical priors, e.g., in the form of simulators, to regularize the training of neural network models. While learning accurate dynamics already in the low data regime, SIM-FSVGD scales and excels also when more data is available. We empirically show that learning with implicit physical priors results in accurate mean model estimation as well as precise uncertainty quantification. We demonstrate the effectiveness of SIM-FSVGD in bridging the sim-to-real gap on a high-performance RC racecar system. Using model-based RL, we demonstrate a highly dynamic parking maneuver with drifting, using less than half the data compared to the state of the art.

## I. INTRODUCTION

For many decades, physical equations of motion have been leveraged to perform highly dynamic and complex tasks in robotics. However, recent advances have shown a significant gap between the physical models and the real world, also known as the ‘sim-to-real’ gap. The gap is primarily because domain-specific models/simulators often tend to neglect complex real-world phenomena (e.g., aerodynamics, system latency, elastic deformation, etc. [1, 2])

Hence, in recent years, data-driven approaches, such as reinforcement learning (RL), have been widely applied in robotics [3]. Particularly, model-based RL methods are often used for robot learning directly on hardware due to their improved sample-efficiency [4–7]. In model-based deep RL, the dynamics of the real-world system are learned directly from data using neural networks (NNs). However, these methods often ignore the abundance of knowledge available in physical simulators and instead directly fit a model from scratch. We believe that this approach is a critical source of data inefficiency.

In this paper, we discuss how we can harness the benefits of both worlds, i.e., having the flexibility of NNs while retaining the domain knowledge embedded in simulators, encapsulating models of various degrees of fidelity. Moreover, we leverage recent advances in Bayesian inference in functional space [8, 9] to impose a functional prior derived from the simulation/physical model. The prior regularizes the network to behave similarly to the simulation in a low-data regime and fit the real-world dynamics when data is available. Empirically, this results in a substantial sample-efficiency gain.

<sup>\*</sup>Equal contribution listed in alphabetical order

<sup>1</sup>Department of Computer Science, ETH Zürich Switzerland email: `firstname.lastname@inf.ethz.ch`

<sup>2</sup>Department of Electrical Engineering, ETH Zürich Switzerland email: `dorfler@ethz.ch`



Fig. 1: Dynamic RC car from our experiments.

In summary, our main contributions are; (i) we propose SIM-FSVGD, a simple and tractable update rule for training Bayesian Neural Networks (BNNs) that incorporates knowledge for available simulators, (ii) we show that SIM-FSVGD has significant gains in sample-efficiency on learning real-world dynamics with NNs over several other widely applied methods, and (iii) we combine SIM-FSVGD with a model-based RL algorithm and show that it results in considerably faster convergence on a highly dynamic RC Car (cf., fig. 1).

## II. RELATED WORK.

1) *Bridging the sim-to-real gap*: Methods such as domain randomization, meta-RL, and system identification are typically used to bridge the sim-to-real gap in robotics and RL (cf., [10]). Domain randomization based approaches assume that the real world is captured well within the randomized simulator. In many cases, such as the ones we consider in this work, this is not the case. Meta-RL methods [11–13, 7] also require that the meta-training and test tasks are i.i.d. samples from a task distribution, and thereby representative. Furthermore, they generally require a plethora of data for meta-training, which is typically also computationally expensive. Instead, SIM-FSVGD does not require any meta training and solely leverages query access to low-fidelity simulators which are based on first-principles physics models and are often available in robotics. Moreover, we leverage solely the real world data and operate in the data-scarce regime. To this end, we consider model-based learning approaches due to their sample efficiency. Typical model-based approaches, perform system identification [14–19] to fit the simulation parameters on real-world data. However, if the simulation model cannot capture the real-world dynamics well, this results in suboptimal performance (cf., section V). Most closely related to our work are approaches such as [2, 20–23]. They use the simulator model and fit an additional model, e.g., NN or Gaussian process, on top to close the sim-to-real gap. We refer to this approach as GREYBOX. In section V we show that our method results in better uncertainty estimates than the GREYBOX approach.

2) *Neural Networks for System Identification*: Neural networks are commonly used for learning robot dynamics

from data [2, 24–28] due to their high expressiveness. However, during control, planning with NN models is challenging because most planners exploit (“overfit” to) inaccuracies of the learned model. This leads to suboptimal performance on the real system [29]. BNNs do not suffer from the same pitfall [29]. As a result, recently, BNNs have been widely used for dynamics learning [29–34].

3) *Bayesian Neural Networks*: Unlike learning a single NN, BNNs maintain a distribution over the NN parameters. Moreover, using a known prior distribution over the NN parameters, given the training data, BNNs infer the posterior distribution. Exact Bayesian inference is computationally intractable for NNs, and approximations [35–41] are generally used. In this work, we focus on particle-based BNNs due to their computational tractability and reliable uncertainty estimates [8, 42, 43] (cf., [44] or section III for more details). Nearly all BNNs impose a generic standard normal prior on the NN parameters. On the contrary, we leverage existing (low-fidelity) simulators to impose an informed prior. This boosts sample efficiency when learning robot dynamics from data and also enables bridging the ‘sim-to-real’ gap (cf., section V).

### III. BACKGROUND

#### A. Learning Dynamics Models with BNNs

The main goal of this work is to efficiently learn a robot dynamics model from data, that is, given a dataset of state-action<sup>1</sup> pairs and the next state  $(s_t, u_t, s_{t+1})$ , learn a model to predict the next state. To this end, we focus on a general regression problem with a dataset  $\mathcal{D} = (\mathbf{X}^{\mathcal{D}}, \mathbf{y}^{\mathcal{D}})$ , comprised of  $m$  noisy evaluations  $y_j = f(\mathbf{x}_j) + \epsilon_j$  of an unknown function  $f: \mathcal{X} \mapsto \mathcal{Y}$ . Here, the training inputs are denoted by  $\mathbf{X}^{\mathcal{D}} = \{\mathbf{x}_j\}_{j=1}^m$  and corresponding function values by  $\mathbf{y}^{\mathcal{D}} = \{y_j\}_{j=1}^m$ . The noise is assumed to be i.i.d. and Gaussian. To fit a regressor on the data, we employ a NN  $h_{\theta}: \mathcal{X} \rightarrow \mathcal{Y}$  with weights  $\theta \in \Theta$ . The conditional predictive distribution for the noisy observations can be defined as  $p(y|\mathbf{x}, \theta) = \mathcal{N}(y|h_{\theta}(\mathbf{x}), \sigma^2)$  using  $h_{\theta}$ , with  $\sigma^2$  as the observation noise variance.

BNNs maintain a distribution over the NN parameters  $\theta$ . In particular, Bayes’ theorem combines a prior distribution  $p(\theta)$  with the empirical data into a posterior distribution  $p(\theta|\mathbf{X}^{\mathcal{D}}, \mathbf{y}^{\mathcal{D}}) \propto p(\mathbf{y}^{\mathcal{D}}|\mathbf{X}^{\mathcal{D}}, \theta)p(\theta)$ . Here,  $p(\mathbf{y}^{\mathcal{D}}|\mathbf{X}^{\mathcal{D}}, \theta)$  is the likelihood of the data, given a NN hypothesis  $\theta$ . Under the i.i.d. hypothesis, the likelihood factorizes as  $p(\mathbf{y}^{\mathcal{D}}|\mathbf{X}^{\mathcal{D}}, \theta) = \prod_{j=1}^m p(y_j|\mathbf{x}_j, \theta)$ .

To make predictions for an unseen test point  $\mathbf{x}^*$ , we marginalize out the parameters  $\theta$  in the posterior, i.e., the predictive distribution is calculated as

$$\begin{aligned} p(y^*|\mathbf{x}^*, \mathbf{X}^{\mathcal{D}}, \mathbf{y}^{\mathcal{D}}) &= \int p(y^*|\mathbf{x}^*, \theta)p(\theta|\mathbf{X}^{\mathcal{D}}, \mathbf{y}^{\mathcal{D}})d\theta \\ &= \mathbb{E}_{\theta} [p(y^*|\mathbf{x}^*, \theta)|\mathbf{X}^{\mathcal{D}}, \mathbf{y}^{\mathcal{D}}]. \end{aligned}$$

<sup>1</sup>Analogous to model-based learning literature, we focus on the fully observable setting.

#### B. BNN Inference in the Function Space

Due to the high-dimensional parameter space  $\Theta$  and the over-parameterized nature of the NN mapping  $h_{\theta}(\mathbf{x})$ , posterior inference for BNNs is challenging. Hence, we need to fall back on some form of *approximate inference* [45, 46] to make BNN inference tractable. Moreover, the complex parametrization also makes choosing an appropriate prior distribution  $p(\theta)$  very hard.

Aiming to alleviate these issues, an alternative approach views BNN inference in the function space, i.e., instead of imposing a prior in the parameter space  $\Theta$ , a prior in the space of regression functions  $h: \mathcal{X} \mapsto \mathcal{Y}$  is used. This results in the posterior  $p(h|\mathbf{X}^{\mathcal{D}}, \mathbf{y}^{\mathcal{D}}) \propto p(\mathbf{y}^{\mathcal{D}}|\mathbf{X}^{\mathcal{D}}, h)p(h)$  [8, 9]. Here,  $p(h)$  is a *stochastic process* prior, e.g., Gaussian process (GP) [47] with index set  $\mathcal{X}$ , taking values in  $\mathcal{Y}$ .

Stochastic processes can be understood as infinite-dimensional random vectors. Hence, they are computationally intractable. However, given finite *measurement sets*  $\mathbf{X} := [\mathbf{x}_1, \dots, \mathbf{x}_k] \in \mathcal{X}^k, k \in \mathbb{N}$ , the stochastic process can be characterized by its marginal distributions of function values  $\rho(\mathbf{h}^{\mathbf{X}}) := \rho(h(\mathbf{x}_1), \dots, h(\mathbf{x}_k))$  (cf. Kolmogorov Extension Theorem, [48]). This allows us to break down functional BNN inference into a more tractable form by re-phrasing it in terms of posterior marginals for measurement sets  $\mathbf{X}$ :  $p(\mathbf{h}^{\mathbf{X}}|\mathbf{X}, \mathbf{X}^{\mathcal{D}}, \mathbf{y}^{\mathcal{D}}) \propto p(\mathbf{y}^{\mathcal{D}}|\mathbf{h}^{\mathbf{X}^{\mathcal{D}}})p(\mathbf{h}^{\mathbf{X}})$ .

This functional posterior can be tractably approximated for BNN inference [8, 9]. In the following, we briefly describe how this can be done with functional Stein Variational Gradient Descent [FSVGD, 8]: The FSVGD method approximates the posterior as a set of  $L$  NN parameter particles  $\{\theta_1, \dots, \theta_L\}$ . To improve the particle approximation, FSVGD iteratively re-samples measurement sets  $\mathbf{X}$  from a distribution  $\nu$  supported on  $\mathcal{X}$ , e.g., UNIFORM( $\mathcal{X}$ ). It then computes SVGD updates [43] in the function space, and projects them back into the parameter space, to update the particle configuration.

To guide the particles towards areas of high posterior probability, FSVGD uses the functional posterior score, i.e., the gradient of the log-density

$$\begin{aligned} \nabla_{\mathbf{h}^{\mathbf{X}}} \ln p(\mathbf{h}^{\mathbf{X}}|\mathbf{X}, \mathbf{X}^{\mathcal{D}}, \mathbf{y}^{\mathcal{D}}) &= \\ \nabla_{\mathbf{h}^{\mathbf{X}}} \ln p(\mathbf{y}^{\mathcal{D}}|\mathbf{h}^{\mathbf{X}^{\mathcal{D}}}) + \nabla_{\mathbf{h}^{\mathbf{X}}} \ln p(\mathbf{h}^{\mathbf{X}}). \end{aligned} \quad (1)$$

In particular, for all  $l = 1, \dots, L$  the particle updates are computed as

$$\theta^l \leftarrow \theta^l - \gamma J_l u_l, \quad \text{where } J_l = \underbrace{(\nabla_{\theta^l} \mathbf{h}_{\theta^l}^{\mathbf{X}})^{\top}}_{\text{NN Jacobian}} \quad (2)$$

$$u_l = \underbrace{\left( \frac{1}{L} \sum_{i=1}^L \mathbf{K}_{li} \nabla_{\mathbf{h}_{\theta^i}^{\mathbf{X}}} \ln p(\mathbf{h}_{\theta^i}^{\mathbf{X}}|\mathbf{X}, \mathbf{X}^{\mathcal{D}}, \mathbf{y}^{\mathcal{D}}) + \nabla_{\mathbf{h}_{\theta^l}^{\mathbf{X}}} \mathbf{K}_{li} \right)}_{\text{SVGD update in the function space}},$$

where  $\mathbf{K} = [k(\mathbf{h}_{\theta^i}^{\mathbf{X}}, \mathbf{h}_{\theta^j}^{\mathbf{X}})]_{li}$  is the kernel matrix between the function values in the measurement points based on a kernel function  $k(\cdot, \cdot)$ . As we can see in (2), FSVGD only uses the prior scores and, in principle, does not require the density

function of the prior marginals. This constitutes a key insight that we will later draw upon in our approach.

#### IV. SIM-FSVGD

In this section, we present SIM-FSVGD. We start this section with an illustrative example to motivate our approach.

##### A. Illustrative Example: the Pendulum

Let us consider a simplified pendulum model where the pendulum’s state  $\mathbf{s} = [\theta, \dot{\theta}]$  is comprised of the angle  $\theta$  and angular velocity  $\dot{\theta}$ . We can describe the change of state (i.e., transition function) by the ODE

$$\ddot{\theta} = \frac{mgl \sin(\theta) + C_m u}{I}. \quad (3)$$

Here  $m, l$  and  $I$  are the mass, length, and moment of inertia, respectively. The motor at the rotational joint of the pendulum applies the torque  $\tau = C_m u$  proportional to the control input  $u$ .

When aiming to accurately predict how the real pendulum system behaves, there are two key sources of uncertainty/error. First, the exact parameters  $\phi$  (here,  $\phi = [m, l, C_m, I]$ ) of the domain-specific model are unknown. However, we can typically narrow down each parameter’s value to a plausible range. This can be captured via a prior distribution  $p(\phi)$  over the model parameters. The process of randomly sampling a parameter set  $\phi \sim p(\phi)$  and then integrating/solving the ODE in Equation (3) with the corresponding parameters gives random functions. This allows us to implicitly construct a stochastic process of functions that reflect our simplified pendulum model. We call this the *domain-model process*.

The second source of uncertainty is various physical phenomena such as aerodynamic drag, friction, and motor dynamics that are not captured in (3). This results in a systematic sim-to-real gap. When constructing a prior  $p(h)$  for training a BNN model on the real system, we also need to model the sim-to-real gap.

##### B. Main Algorithm

Now we describe our approach, SIM-FSVGD, which combines the domain-model process with a sim-to-real prior to construct an informed functional prior  $p(h)$  for learning the dynamics model. The final algorithm is presented in Algorithm 1. First, we summarize how to construct the stochastic process prior.

In practice, we aim to model vector-valued functions, i.e.,  $h : \mathcal{X} \mapsto \mathcal{Y}$  where  $\mathcal{X} \subseteq \mathbb{R}^{d_x}$  and  $\mathcal{Y} \subseteq \mathbb{R}^{d_y}$ . We factorize the prior over the output dimensions, i.e.,  $p(h) = \prod_{i=1}^{d_y} p(h_i)$ . This allows us to treat each  $h_i : \mathbf{X} \mapsto \mathbb{R}$  as an independent scalar-valued function.

1) *Domain-model process*: The first component of the implicit prior is the domain-specific (low-fidelity) simulation model  $g(\mathbf{x}, \phi)$  which takes as an input the state action pair  $(\mathbf{s}, \mathbf{u})$ , i.e.,  $\mathbf{x} = [\mathbf{s}, \mathbf{u}]$ , as well as system parameters  $\phi$ . Crucially, we do not require an analytical expression of  $g : \mathcal{X} \mapsto \mathcal{Y}$ . It is sufficient to have query access, i.e., the

output vector of  $g(\mathbf{x}, \phi)$  can be the result of a numerical simulation (e.g., ODE solver or rigid body simulator).

2) *Sim-to-real prior*: The second component is the sim-to-real gap process. For that we employ a GP  $p(\tilde{h}_i)$  per output dimension  $i = 1, \dots, d_y$  which aims to model the residual/gap between the best possible domain model  $g_i(\mathbf{x}, \phi^*)$  and the actual target function  $f(\mathbf{x})$ . We use a GP with zero mean and an isotropic kernel  $k(\mathbf{x}, \mathbf{x}') = \nu^2 \rho(\|\mathbf{x} - \mathbf{x}'\|/\ell)$  with variance  $\nu^2$  and lengthscale  $\ell$ . Accordingly, the marginal distributions follow a multivariate normal distribution  $p(\tilde{\mathbf{h}}_i^{\mathbf{X}}) = \mathcal{N}(\tilde{\mathbf{h}}_i^{\mathbf{X}} | \mathbf{0}, \mathbf{K})$  with kernel matrix  $\mathbf{K} = [k(\mathbf{x}_i, \mathbf{x}_j)]_{i,j}$ . The lengthscale  $\ell$  and  $\nu^2$  are hyperparameters that influence our belief on the sim-to-real gap. A small  $\nu^2$  implies that the sim-to-real gap is with high probability not too big and a large lengthscale conveys that deviations from the domain model are systematic rather than local.

We combine both processes by independently sampling (conditional) random vectors from both processes and adding them, i.e.,

$$\mathbf{h}_i^{\mathbf{X}} = [g_i(\mathbf{x}_1, \phi), \dots, g_i(\mathbf{x}_k, \phi)]^\top + \tilde{\mathbf{h}}_i^{\mathbf{X}} \quad (4)$$

with  $\phi \sim p(\phi)$ ,  $\tilde{\mathbf{h}}_i^{\mathbf{X}} \sim \mathcal{N}(\tilde{\mathbf{h}}_i^{\mathbf{X}} | \mathbf{0}, \mathbf{K})$ .

The resulting stochastic process prior  $p(h)$  is defined implicitly through the marginal distributions that are implied by (4).

3) *Estimating the Stochastic Process Score*: To perform FSVGD updates as in (2), we have to compute the score of our stochastic process prior, i.e.,  $\nabla_{\mathbf{h}^{\mathbf{X}}} \ln p(\mathbf{h}^{\mathbf{X}}) = \sum_i^{d_y} \nabla_{\mathbf{h}_i^{\mathbf{X}}} \ln p(\mathbf{h}_i^{\mathbf{X}})$ . This is intractable for most domain-model processes and we need to estimate it from samples. To this end, we use a simple and efficient Gaussian approximation for the score of prior marginals. In particular, we sample a measurement set  $\mathbf{X}$  from a known measurement distribution  $\nu$ , e.g., uniform distribution over the state-action space. For the measurement set  $\mathbf{X}$ , we sample  $P$  vectors of function values  $\mathbf{h}_{i,1}^{\mathbf{X}}, \dots, \mathbf{h}_{i,P}^{\mathbf{X}} \sim p(\mathbf{h}_i^{\mathbf{X}})$  via (4). Then, we compute their mean and covariance matrix, that is  $\mu_i^{\mathbf{X}}$  and  $\Sigma_i^{\mathbf{X}}$  and use the score of the corresponding multivariate Gaussian as an approximation, i.e.,  $p(\mathbf{h}_i^{\mathbf{X}}) \sim \mathcal{N}(\mu_i^{\mathbf{X}}, \Sigma_i^{\mathbf{X}})$ .

#### V. EXPERIMENTS

In this section, we present our experimental results where we evaluate our method, SIM-FSVGD, on learning the robot dynamics from data and online reinforcement learning.

a) *Baselines*: We consider three baselines; (i) FSVGD [8], (ii) SYSID, and (iii) GREYBOX. As discussed in section III, FSVGD is widely applied for Bayesian inference in deep learning. SYSID is a simple baseline, that uses the data from the real system to fit the parameters of the simulation model. Crucially, SYSID cannot capture the sim-to-real gap. To this end, we also consider GREYBOX which first fits the simulation model and then learns a BNN to fit the difference between the real data and the sim model, akin to [20, 22, 23].

##### A. One-Dimensional Experiment

We consider a one-dimensional sinusoidal function to visualize the difference between our method and the baselines

---

## Estimating Functional Prior from Physical Prior

---

**Input:** Measurement distribution  $\nu$ , Physical prior  $g$ , Parameter distribution  $p(\phi)$ .

- 1: Sample measurement set  $\mathbf{X}_M \sim \nu$
  - 2: Sample parameters  $\phi_j$  for  $j \in \{1, \dots, P\}$  from  $p(\phi)$ .
  - 3: Calculate physical model output  $g(\mathbf{X}_M; \phi_j)$  for  $j \in \{1, \dots, P\}$ .
  - 4: Calculate empirical mean  $\mu(\mathbf{X}_M)$  and variance  $\Sigma(\mathbf{X}_M)$  over the  $P$  samples.
  - 5: Approximate Gaussian prior  $p(\mathbf{h}^{\mathbf{X}_M}) \sim \mathcal{N}(\mu(\mathbf{X}_M), \Sigma(\mathbf{X}_M))$
- 




Fig. 2: BNN posteriors trained on two data points for a one-dimensional sinusoidal function. GREYBOX finds accurate mean, but has large uncertainty estimates. At the same time SIM-FSVGD obtains both accurate mean and uncertainty estimates.

---

### Algorithm 1 SIM-FSVGD

---

**Input:** Measurement distribution  $\nu$ , Sim prior  $g$ , Parameter distribution  $p(\phi)$ , GP  $p(\mathbf{h})$ , Data  $\mathcal{D}$ , BNN particles  $\{\theta_i\}_{i=1}^L$

- 1: Sample measurement set  $\mathbf{X} \sim \nu$
  - 2: Sample sim functions  $\{\mathbf{h}_{i,j}^{\mathbf{X}}\}_{i=0}^{d_y}$  for  $j \in \{1, \dots, P\}$ .
  - 3: Calculate empirical mean  $\{\mu_i^{\mathbf{X}}\}_{i=1}^{d_y}$  and variance  $\{\Sigma_i^{\mathbf{X}}\}_{i=0}^{d_y}$  over the  $P$  samples.
  - 4: Approximate Gaussian prior for each  $i$ , i.e.  $p(\mathbf{h}_i^{\mathbf{X}}) \sim \mathcal{N}(\mu_i^{\mathbf{X}}, \Sigma_i^{\mathbf{X}})$
  - 5: Update BNN particles with eq. (2).
- 

in fig. 2. Moreover, we train the BNNs with only two data points and plot their posteriors. From fig. 2 we conclude that FSVGD and GREYBOX result in very general posteriors that do not capture the sinusoidal nature of the true function. On the contrary, SIM-FSVGD can learn an accurate posterior with only two data points.

### B. Sim to Real Transfer

In this experiment, we investigate how SIM-FSVGD bridges the sim-to-real gap for dynamical systems. To this end, we consider a pendulum and racecar in simulation as well as a real-world RC car as our systems. We use an i.i.d. dataset of sampled transitions from the real system to fit a model of the dynamics. As the evaluation metric, we consider the negative log-likelihood (NLL) of our model on a test dataset. For the pendulum, we use the model from eq. (3) as the simulation prior and we simulate the sim-to-real gap by incorporating aerodynamic drag, friction, and motor dynamics in the ‘real’ dynamics. For the racecar, we use the dynamics model from [49] which models tire dynamics with the

Pacejka tire model [50]. For the simulation prior, we use the kinematics bicycle model also from [49], which does not capture the tire dynamics. We use the same simulation prior for real RC car. The car consists of a high-torque motor, which allows us to perform dynamic maneuvers that involve loss of traction and drifting. These behaviors are neither captured by our simulation prior nor the dynamics model from [49].

The RC car is similar to the one in [28, 7] and has a six-dimensional state (position, orientation, and velocities) and two-dimensional input (steering and throttle). We control the car at 30 Hz. and use the Optitrack for robotics motion capture system<sup>2</sup> to estimate the state. The car has a significant delay (ca. 80 ms) between transmission and execution of the control signals on the car. Therefore, we include the last three actions  $[a_{t-3}, a_{t-2}, a_{t-1}]$  in the current state  $s_t$ . The resulting state space is  $12D$  and the action space is  $2D$ .

Our results are presented in fig. 3. From the figure, we conclude that SIM-FSVGD consistently outperforms all other methods in both case studies. Moreover, in the low-data regime, our method achieves orders of magnitude lower NLL, even though there is a considerable sim-to-real gap (c.f., SYSID). Furthermore, when more data is provided our method performs similarly to other baselines. Therefore, SIM-FSVGD relies less on the simulation prior when more data is available. This empirically validates the consistency of our approach. In fig. 3, we also study the effect of incorporating the additional sim-to-real GP prior in SIM-FSVGD. We conclude that including the additional GP leads to better performance and thus plays a crucial role, particularly for the pendulum and racecar.




Fig. 3: We benchmark NLL of dynamics model on test data for Racecar on hardware and in simulations and Pendulum in simulation. BNNs with simulation priors achieve low NLL already in small data regimes and constantly outperform the standard BNNs when we increase the number of train data points.




Fig. 4: Desired reverse parking maneuver which involves rotating the car 180° and parking ca. 2 m away.




Fig. 5: Realized trajectory for Racecar environment in simulation after episode 3. In violet, we have the trajectory of the model with FSVGD (no prior), and in red the trajectory of SIM-FSVGD using the low-fidelity prior. We observe that, SIM-FSVGD already learns to get close to the target in three episodes, whereas FSVGD (no prior) explores away from it.

### C. Online RL

We study the problem of model-based reinforcement learning for the racecar in simulation and the real RC car. In model-based RL, BNNs with good uncertainty estimates play a crucial role in exploration [29–33]. Moreover, incorporating a simulation prior can lead to directed exploration and therefore faster convergence, i.e., better sample efficiency. To this end, in this experiment, we compare the performance of SIM-FSVGD and FSVGD in model-based RL. Moreover, we use our BNN to obtain a policy  $\pi$ , which we rollout on the real system to get further data for updating our model (see algorithm 2). We repeat this procedure for 20 episodes and consider a rollout length of 100. For the policy  $\pi$ , we use the SAC algorithm [51] which we train using our BNN similar to [52]. The RL task we consider involves

### Algorithm 2 SIM-FSVGD for Model-Based RL

---

**Input:** Initial model and policy  $f_{\Theta}, \pi_{\psi}, \mathcal{D} := \emptyset$

- 1: **for** episode = 1, 2, ... **do**
- 2:    $f_{\Theta} \leftarrow \text{SIM-FSVGD}(f_{\Theta}, \mathcal{D})$
- 3:    $\pi_{\psi} \leftarrow \text{SAC}(\pi_{\psi}, f_{\Theta})$
- 4:    $(s_0, a_0, \dots, a_{T-1}, s_T) \leftarrow \text{ROLLOUT}(\pi_{\psi})$
- 5:    $\mathcal{D} \leftarrow \mathcal{D} \cup \{(s_t, a_t, s_{t+1})\}_{t=0}^{T-1}$
- 6: **end for**

---

reverse parking the car ca. 2 m away from the initial position  $x_0$  by quickly rotating the car 180°. (c.f., fig. 4 or the video demonstration<sup>3</sup>). This typically leads to the car drifting, a behavior that is not captured by our simulation prior.

In Figure 7 we report the learning curves for the racecar in simulation and the real RC car. We observe that in both cases, incorporating the simulation prior leads to considerably faster convergence (ca. 2× fewer episodes). Moreover, in Figure 3 we show the realized trajectory for the third episode of the racecar. We see that already after three episodes, SIM-FSVGD learns to drive the car close to the target whereas not incorporating the prior leads to exploration away from the target position. Similar behavior is observed on hardware (c.f., Figure 6), where from episode 1, the policy learned using SIM-FSVGD drives the car close to the target, and at episode 10 it is parking the car nearly perfectly. On the contrary, not incorporating the prior leads to no movement of the car in episode 1 and the car driving away from the target in episode 10. This experiment demonstrates that incorporating a simulation prior in model-based RL leads to directed exploration which results in better sample-efficiency.

<sup>2</sup><https://optitrack.com/applications/robotics/>

<sup>3</sup>[https://crl.ethz.ch/videos/iros\\_video](https://crl.ethz.ch/videos/iros_video)




Fig. 6: Realized trajectories after episodes one and ten for SIM-FSVGD and FSVGD. SIM-FSVGD moves close to the target in episode one and parks nearly perfectly in episode ten. On the contrary, FSVGD fails to move at all in the first episode and explores away from the target in the tenth.




Fig. 7: The simulation prior helps initially by exploring more directly towards the target. Both in simulation and on hardware the agent solves the task significantly faster and with less collected data.

## VI. CONCLUSION

In this work, we address the sim-to-real gap in robotics by harnessing both domain knowledge and data. Leveraging ideas from Bayesian inference in functional space, we propose SIM-FSVGD, a simple and tractable algorithm to incorporate physical simulation priors in learning neural network robot dynamics. Our resulting algorithm benefits from both the expressiveness of neural networks and the domain knowledge from physical simulators. We show that SIM-FSVGD results in significantly better sample efficiency for supervised learning of robot dynamics and model-based RL.

**Future Work.** Here, we use a simple Gaussian approximation for score estimation. There are a range of other score estimation techniques [e.g. 53, 54]. We leave a detailed comparison among the different methods as future work. Moreover, we regard applying SIM-FSVGD to more high-dimensional systems from robotics or other scientific domains as an interesting avenue for future work.

## VII. ACKNOWLEDGMENTS

This research was supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program grant agreement no. 815943, the Swiss National Science Foundation under NCCR Automation, grant agreement 51NF40 180545, and the Mi-

crosoft Swiss Joint Research Center. Jonas Rothfuss was supported by an Apple Scholars in AI/ML fellowship.

## REFERENCES

- [1] D. Widmer, D. Kang, B. Sukhija, J. Hübotter, A. Krause, and S. Coros, “Tuning legged locomotion controllers via safe bayesian optimization,” *CORL*, 2023.
- [2] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, and M. H. Vladlen Koltun, “Learning agile and dynamic motor skills for legged robots,” *Science Robotics*, 2019.
- [3] B. Singh, R. Kumar, and V. P. Singh, “Reinforcement learning in robotic applications: a comprehensive survey,” *AI Review*, 2022.
- [4] M. Deisenroth and C. Rasmussen, “PILCO: A Model-Based and Data-Efficient Approach to Policy Search.”
- [5] A. S. Polydoros and L. Nalpantidis, “Survey of model-based reinforcement learning: Applications on robotics,” *Journal of Intelligent & Robotic Systems*, 2017.
- [6] P. Wu, A. Escontrela, D. Hafner, P. Abbeel, and K. Goldberg, “Daydreamer: World models for physical robot learning,” in *CORL*, 2023.
- [7] A. Bhardwaj, J. Rothfuss, B. Sukhija, Y. As, M. Hutter, S. Coros, and A. Krause, “Data-efficient task generalization via probabilistic model-based meta reinforcement learning,” *arXiv*, 2023.
- [8] Z. Wang, T. Ren, J. Zhu, and B. Zhang, “Function space particle optimization for Bayesian neural networks,” in *ICLR*, 2019.
- [9] S. Sun, G. Zhang, J. Shi, and R. Grosse, “Functional variational bayesian neural networks,” in *ICLR*, 2019.
- [10] W. Zhao, J. P. Queralta, and T. Westerlund, “Sim-to-real transfer in deep reinforcement learning for robotics: a survey,” in *IEEE SSCI*, 2020.
- [11] Y. Duan, J. Schulman, X. Chen, P. L. Bartlett, I. Sutskever, and P. Abbeel, “RI2: Fast reinforcement

- learning via slow reinforcement learning,” in *ICLR*, 2017.
- [12] C. Finn, P. Abbeel, and S. Levine, “Model-agnostic meta-learning for fast adaptation of deep networks,” in *ICML*, 2017.
- [13] C. Finn, K. Xu, and S. Levine, “Probabilistic model-agnostic meta-learning,” in *NeurIPS*, 2018.
- [14] L. Ljung, *System Identification*. John Wiley & Sons, Ltd, 1999.
- [15] C. Li, T. Zhang, and D. I. Goldman, “A terradynamics of legged locomotion on granular media,” *science*, 2013.
- [16] R. Moeckel, Y. N. Perov, A. T. Nguyen, M. Vespignani, S. Bonardi, S. Pouya, A. Sproewitz, J. van den Kieboom, F. Wilhelm, and A. J. Ijspeert, “Gait optimization for roombots modular robots—matching simulation and reality,” in *IROS*, 2013, pp. 3265–3272.
- [17] J. Tan, Z. Xie, B. Boots, and C. K. Liu, “Simulation-based design of dynamic controllers for humanoid balancing,” in *IROS*. IEEE, 2016, pp. 2729–2736.
- [18] S. Zhu, A. Kimmel, K. E. Bekris, and A. Boularias, “Model identification via physics engines for improved policy search,” *IJCAI*, 2017.
- [19] S. Zimmermann, R. Poranne, and S. Coros, “Go fetch! - dynamic grasps using boston dynamics spot with external robotic arm,” in *ICRA*, 2021.
- [20] P. Pastor, M. Kalakrishnan, J. Binney, J. Kelly, L. Righetti, G. Sukhatme, and S. Schaal, “Learning task error models for manipulation,” in *ICRA*, 2013.
- [21] S. Ha and K. Yamane, “Reducing hardware experiments for model learning and policy optimization,” in *ICRA*, 2015.
- [22] H. Huang, W. He, and S. Zhang, “Neural network control of a robotic manipulator with time-varying output constraints by state feedback,” *IFAC*, 2017.
- [23] A. Schperberg, Y. Tanaka, F. Xu, M. Menner, and D. Hong, “Real-to-sim: Predicting residual errors of robotic systems with sparse data using a learning-based unscented kalman filter,” in *UR*, 2023.
- [24] K. Narendra and K. Parthasarathy, “Identification and control of dynamical systems using neural networks,” *IEEE Transactions on Neural Networks*, 1990.
- [25] J. Sjöberg, H. Hjalmarsson, and L. Ljung, “Neural networks in system identification,” *IFAC Proceedings Volumes*, 1994.
- [26] A. Nagabandi, G. Kahn, R. S. Fearing, and S. Levine, “Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning,” in *ICRA*, 2018.
- [27] J. M. Bern, Y. Schneider, P. Banzet, N. Kumar, and S. Coros, “Soft Robot Control With a Learned Differentiable Model,” in *RoboSoft*, 2020.
- [28] B. Sukhija, N. Köhler, M. Zamora, S. Zimmermann, S. Curi, A. Krause, and S. Coros, “Gradient-based trajectory optimization with learned dynamics,” in *ICRA*, 2023.
- [29] K. Chua, R. Calandra, R. McAllister, and S. Levine, “Deep Reinforcement Learning in a Handful of Trials using Probabilistic Dynamics Models,” in *NeurIPS*, 2018.
- [30] S. Curi, F. Berkenkamp, and A. Krause, “Efficient model-based reinforcement learning through optimistic policy search and planning,” *NeurIPS*, 2020.
- [31] R. Sekar, O. Rybkin, K. Daniilidis, P. Abbeel, D. Hafner, and D. Pathak, “Planning to explore via self-supervised world models,” in *ICLR*, 2020.
- [32] L. Treven, J. Hübotter, B. Sukhija, F. Dörfler, and A. Krause, “Efficient exploration in continuous-time model-based reinforcement learning,” *NeurIPS*, 2024.
- [33] B. Sukhija, L. Treven, C. Sancaktar, S. Blaes, S. Coros, and A. Krause, “Optimistic active exploration of dynamical systems,” *NeurIPS*, 2024.
- [34] J. Rothfuss, B. Sukhija, T. Birchler, P. Kassraie, and A. Krause, “Hallucinated adversarial control for conservative offline policy evaluation,” *UAI*, 2023.
- [35] D. J. MacKay, “Bayesian methods for adaptive models,” Ph.D. dissertation, California Institute of Technology, 1992.
- [36] R. M. Neal, *Bayesian learning for neural networks*. Springer Science & Business Media, 2012.
- [37] A. Graves, “Practical variational inference for neural networks,” *NeurIPS*, 2011.
- [38] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, “Weight uncertainty in neural network,” in *ICML*, 2015.
- [39] Y. Gal and Z. Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning,” in *ICML*, 2016.
- [40] R. Zellers, Y. Bisk, R. Schwartz, and Y. Choi, “Swag: A large-scale adversarial dataset for grounded common-sense inference,” *arXiv*, 2018.
- [41] E. Goan and C. Fookes, “Bayesian neural networks: An introduction and survey,” *Case Studies in Applied Bayesian Data Science*, 2020.
- [42] B. Lakshminarayanan, A. Pritzel, and C. Blundell, “Simple and scalable predictive uncertainty estimation using deep ensembles,” *NeurIPS*, 2017.
- [43] Q. Liu and D. Wang, “Stein Variational Gradient Descent: A General Purpose Bayesian Inference Algorithm,” in *NeurIPS*, 2016.
- [44] F. D’Angelo and V. Fortuin, “Repulsive deep ensembles are bayesian,” *NeurIPS*, 2021.
- [45] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, “Variational inference: A review for statisticians,” *Journal of the American Statistical Association*, 2017.
- [46] C. Chen, R. Zhang, W. Wang, B. Li, and L. Chen, “A unified particle-optimization framework for scalable bayesian sampling,” 2018.
- [47] C. E. Rasmussen and C. K. I. Williams, *Gaussian processes in machine learning*. MIT Press, 2006.
- [48] B. Øksendal, *Stochastic Differential Equations: An Introduction with Applications*, 2000.
- [49] J. Kabzan, M. I. Valls, V. J. Reijgwart, H. F. Hendriks, C. Ehmke, M. Prajapat, A. Bühler, N. Gos-

ala, M. Gupta, R. Sivanesan, *et al.*, “Amz driverless: The full autonomous racing system,” *Journal of Field Robotics*, 2020.

- [50] H. B. Pacejka and E. Bakker, “The magic formula tyre model,” *Vehicle system dynamics*, pp. 1–18, 1992.
- [51] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” in *ICML*, 2018.
- [52] M. Janner, J. Fu, M. Zhang, and S. Levine, “When to trust your model: Model-based policy optimization,” *NeurIPS*, 2019.
- [53] J. Shi, S. Sun, and J. Zhu, “A spectral approach to gradient estimation for implicit distributions,” in *ICML*, 2018.
- [54] Y. Zhou, J. Shi, and J. Zhu, “Nonparametric score estimators,” in *ICML*, 2020.