

WHICH MODEL TO TRUST: ASSESSING THE INFLUENCE OF MODELS ON THE PERFORMANCE OF REINFORCEMENT LEARNING ALGORITHMS FOR CONTINUOUS CONTROL TASKS

Anonymous authors

Paper under double-blind review

ABSTRACT

The need for algorithms able to solve Reinforcement Learning (RL) problems with few trials has motivated the advent of model-based RL methods. The reported performance of model-based algorithms has dramatically increased within recent years. However, it is not clear how much of the recent progress is due to improved algorithms or due to improved models. While different modeling options are available to choose from when applying a model-based approach, the distinguishing traits and particular strengths of different models are not clear. The main contribution of this work lies precisely in assessing the model influence on the performance of RL algorithms. A set of commonly adopted models is established for the purpose of model comparison. These include Neural Networks (NNs), ensembles of NNs, two different approximations of Bayesian NNs (BNNs), that is, the Concrete Dropout NN and the Anchored Ensembling, and Gaussian Processes (GPs). The model comparison is evaluated on a suite of continuous control benchmarking tasks. Our results reveal that significant differences in model performance do exist. The Concrete Dropout NN reports persistently superior performance. We summarize these differences for the benefit of the modeler and suggest that the model choice is tailored to the standards required by each specific application.¹

1 INTRODUCTION

Reinforcement Learning (RL) aims to learn an optimal strategy in a sequential decision-making setting through the interactions between an agent and their environment (Li, 2017). In recent years, Deep RL has been proven capable of handling high-dimensional problems catering to dramatically accelerated progress in the domain of data-driven policy planning, allowing to solve previously intractable problems (Arulkumaran et al., 2017; Mnih et al., 2015; Silver et al., 2017). Nevertheless, the large amount of training data that is required has led to limited adoption in practical real-life settings (Polydoros & Nalpantidis, 2017). As a result, Model-Based Reinforcement Learning (MBRL) has gained increased attention as this approach is expected to close the gap between simulated and real world tasks (Nagabandi et al., 2020), due to the higher sample-efficiency compared to model-free algorithms (Chua et al., 2018; Gal et al., 2016; Janner et al., 2019; Nagabandi et al., 2018; Wang et al., 2019). MBRL is characterized by the utilization of a model, in most cases a Neural Network (NN) or a Gaussian Process (GP), to learn the dynamics of the agent’s environment. The learned model is exploited by an RL algorithm, which attempts to solve the underlying planning task.

While the performance of MBRL algorithms has been increasing dramatically in recent years, when coming to the choice of an appropriate model, the decision remains not trivial. Several authors proposed diverse possibilities for structuring model-based algorithms. Nagabandi et al. (2018) applied a deterministic NN. Deisenroth & Rasmussen (2011) used Gaussian Processes. Later, Gal et al. (2016) improved the previous method by leveraging Bayesian NNs relying on the Monte Carlo (MC) Dropout approximation (Gal & Ghahramani, 2016). Janner et al. (2019) and Chua et al.

¹Code available on GitHub. The link will be published after the review.

(2018) exploited an ensemble of probabilistic NNs. However, it is not common practice to formally investigate the choice of model within MBRL and it is still unclear whether substantial differences exist between available model classes. As a result, it is not clear how much of the recent progress in MBRL is due to improved algorithms or due to improved models. It is therefore evident that a benchmark is needed for providing clarity on the matter.

The main contribution of this work is to establish such a benchmark. We examine a set of mathematical methods that are commonly used as models in MBRL to assess the model influence on RL algorithms. The evaluated models are Neural Networks, ensembles of Neural Networks, Bayesian Neural Networks, and Gaussian Processes. The model comparison is carried out on a suite of 6 different continuous control environments of increasing complexity that are commonly utilized for the performance evaluation of RL algorithms. Since sample-efficiency is the main reason to choose model-based algorithms, we are particularly interested in evaluating how fast a model-based agent can achieve good performance on a specific task. As such, each task is run for a limited number of trials. Besides, we are also interested in overall performance, model robustness, and computing time.

Works that formally assess model influence are lacking in MBRL literature. To the best of our knowledge, Chua et al. (2018) and Wang et al. (2019) represent the only ones that even got close to our intent. Chua et al. (2018) present one of scarce - if not the only - attempt to compare their proposed method with different modeling options. However, the work in Chua et al. (2018) does not aim to provide a benchmark to assess model influence on MBRL given the limited number of tasks and evaluated models. Wang et al. (2019) recognized the difficulty in offering an objective comparison between available MBRL methods. In order to quantify scientific progress, the authors proposed the first benchmark in MBRL. However, their work differs from ours as it focuses on comparing the algorithms rather than the models. While it allowed further steps in understanding MBRL, model choice still remains unclear. In contrast to these prior works, we focus on the aspect of model class selection and aim to assess their influence on the performance of MBRL algorithms.

2 BACKGROUND

In RL problems an agent observes the environment’s state $s_t \in \mathcal{S}$, it selects an action $a_t \in \mathcal{A}$ and receives the reward r_t for the decision made. Next, the environment transitions to the next state $s_{t+1} \in \mathcal{S}$ and a new step of the sequential decision making problem initiates. The goal is to learn a policy, i.e. a sequence of actions over a prescribed horizon, that maximizes the sum of future rewards (Sutton & Barto, 2018).

Formally, a RL problem can be stated as a Markov Decision Process (MDP) (Van Otterlo & Wiering, 2012). In this work, we consider RL problems that can be described as fully observable MDPs defined by the tuple $(\mathcal{S}, \mathcal{A}, f, r, H)$. Here, \mathcal{S} and \mathcal{A} are the state and the action space, respectively, which are both assumed continuous. Function $f : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ is the deterministic transition dynamics $f(s_t, a_t) = s_{t+1}$ of the environment, while $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{R}$ represents the reward function $r(s_t, a_t) = r_t$. In MBRL, this is usually assumed known (Nagabandi et al., 2018; Wang et al., 2019). Finally, H is the horizon of the problem: the considered MDPs are episodic, namely the state is reset after each episode of length H .

In MBRL, the interactions with the environment are collected in a dataset $\mathcal{D} = \{(s_t, a_t, s_{t+1})\}$. The agent thus learns a model of the transition dynamics of the environment $\hat{f}(s_t, a_t) = s_{t+1}$. This is exploited to simulate the environment and predict the next states for a sequence of actions. Hence, the learned model can be used for planning the optimal sequence of actions to be executed in this environment. In order to define such a policy, techniques such as stochastic or trajectory optimization are commonly applied.

Finally, central to MBRL is the concept of uncertainty. Following Kendall & Gal (2017) we distinguish two kinds of uncertainty. Aleatoric or statistical uncertainty is the variance produced by some noise intrinsic in data. Epistemic or model uncertainty is the uncertainty over model parameters and depends on the limited amount of training data available. Modeling the epistemic uncertainty is crucial in MBRL to cope with an imperfect model and know when “the model does not know” (Kendall & Gal, 2017). In addition, it can also drive agent exploration (see Section 4.3).

3 MODELS

This section introduces the models utilized for our evaluation and the motivations that make them interesting modeling choices within MBRL. Furthermore, these models have already been applied successfully to MBRL problems (Nagabandi et al., 2018; Chua et al., 2018; Clavera et al., 2018; Gal et al., 2017; Pearce et al., 2018; 2020; Deisenroth & Rasmussen, 2011; Deisenroth et al., 2013).

3.1 DETERMINISTIC NN

NNs are powerful models, which led to the rise of Deep Reinforcement Learning (Arulkumaran et al., 2017). We refer to this model as fully deterministic: given the model parameters, there is a unique prediction for a specific input. Besides, this model is not able to recover the full distribution that generated the observations nor to model any kind of uncertainty (Chua et al., 2018). As all other NN-based methods, it is also referred to as a model of finite capacity (Gal et al., 2016), which needs to smooth over all data points. Moreover, NN-based models take several training steps for learning, in contrast to GPs. On the other hand, NNs scale very well with dimensionality and are able to handle complex tasks. The Deterministic NN is also the model adopted by the model-based Deep RL pioneering work in Nagabandi et al. (2018) for learning the dynamics of the environment. While many advances have been produced in recent years over this model, some of which are also introduced in the following sections, it is still meaningful to select it for our comparison to draw a baseline and understand whether these constitute actual improvements.

3.2 DETERMINISTIC ENSEMBLE

We define the Deterministic Ensemble as an ensemble of several deterministic NNs. It is well-known that by aggregating a group of predictors, this leads in improved and more robust performance than any single predictor alone (Zhou, 2019). Further, an ensemble of NNs extends the previously described model by providing an estimate of model uncertainty: the variance of the networks' predictions is an approximate measure of the epistemic uncertainty (Beluch et al., 2018). This technique has already demonstrated interesting results (Lakshminarayanan et al., 2016). While it is important that the NNs are mutually independent, it was empirically proved that bootstrapping is unnecessary, and the stochasticity of the optimizer and the random initialization of the weights make the NNs sufficiently independent (Lee et al., 2015). However, the NNs that compose the ensemble are fully deterministic models and do not present any inherent regularization (such as a prior distribution over the weights). As a result, the NNs may converge to similar parameters, nullifying the benefits of the ensemble. Furthermore, deterministic NNs are also more prone to overfit data compared to Bayesian models (Blundell et al., 2015).

3.3 CONCRETE DROPOUT NN

Concrete Dropout (Gal et al., 2017) is a practical approximation for applying the framework of Bayesian deep learning, which promises improved robustness and superior capability for well-calibrated uncertainty estimates. In general, Bayesian NNs are probabilistic models which replace the deterministic parameter values of classical NNs with entire parameter distributions over the weights (MacKay, 1992). Given a prior distribution $p(\theta)$ and some data $D = (\mathbf{x}, \mathbf{y})$, the posterior distribution $p(\theta|D)$ is computed via Bayes rule. Given new input data \mathbf{x}_{new} , it is thus possible to recover the posterior predictive distribution by marginalizing over all posterior parameters:

$$p(\mathbf{y}|\mathbf{x}_{\text{new}}, D) = \int_{\theta} p(\mathbf{y}|\mathbf{x}_{\text{new}}, \theta)p(\theta|D) \quad (1)$$

While it is possible to apply this framework exactly for small-scale problems, computing $p(\theta|D)$ and $p(\mathbf{y}|\mathbf{x}_{\text{new}}, D)$ is generally analytically intractable due to the large number of parameters and to the nonlinear functional form of a NN (Blundell et al., 2015). If we assume $p(\theta|D)$ known, the posterior predictive distribution is empirically recovered with MC approximation. Let $\mathbf{f}^{\hat{\theta}}(\mathbf{x}_i)$ be the output of the BNN for a sample $\hat{\theta} \sim p(\theta|D)$ and a given input \mathbf{x}_i . If one repeats the prediction T times, every time a different sample from the posterior weight distribution is obtained, which yields different outputs. That is, we are taking samples from the posterior predictive distribution

$\mathbf{f}^{\hat{\theta}}(\mathbf{x}_i) \sim p(\mathbf{y}|\mathbf{x}_i, D)$. Provided T is large enough, we can approximate the predictive mean as follows:

$$\mathbb{E}(\mathbf{y}) \approx \frac{1}{T} \sum_{t=1}^T \mathbf{f}^{\hat{\theta}}(\mathbf{x}_i) \quad (2)$$

Unfortunately, recovering $p(\theta|D)$ is an extremely challenging task and the need for frameworks that may sufficiently approximate the posterior weight distribution has long limited the use of BNNs. Concrete Dropout forms one of the two solutions we employ in this work.

Dropout is a technique widely adopted in deep learning for enhancing generalization and improving performance in out-of-sample predictions (Hinton et al., 2012). At each training step, every parameter has a probability p to be turned off and a probability $1 - p$ to be turned on. As a result, a different NN configuration is sampled at each training step. Gal & Ghahramani (2016) have proved that turning on dropout at testing time (called MC dropout) can be interpreted as a variational inference approximation of BNNs with a Bernoulli variational distribution. While MC dropout gained large popularity in supervised learning tasks, two reasons inhibited its application in RL. First, the technique requires hand tuning over the dropout probabilities, with prohibitive computing time in RL applications. Second, the dropout probability should be adapted to the data at hand, making this operation also infeasible within RL where the amount of data changes over time. Ideally, the dropout rate should be higher at the beginning of the RL training and it should decrease as the amount of data increases. To this end, Gal et al. (2017) conceived the Concrete Dropout variant tailored for RL applications. This can be seen as a continuous relaxation of the discrete dropout technique. Consequently, the Concrete Dropout is able to use the derivative estimator to self-optimize the dropout rate for each layer and at each training step, solving both the issues. The Concrete Dropout matches the performance of hand-tuned MC dropout on several tasks and allows the use of BNNs with dropout in RL tasks. Following Gal et al. (2017) and Kendall & Gal (2017), the NN output is decomposed into:

$$\mathbf{f}^{\hat{\theta}}(\mathbf{x}) = (\hat{\mathbf{y}}, \hat{\sigma}) \quad (3)$$

The predictive mean distribution is then computed as in Equation 2, while the predictive variance is approximated as:

$$\text{Var}(\mathbf{y}) \approx \left[\frac{1}{T} \sum_{i=1}^T \hat{\mathbf{y}}_i^2 - \left(\frac{1}{T} \sum_{i=1}^T \hat{\mathbf{y}}_i \right)^2 \right] + \frac{1}{T} \sum_{i=1}^T \hat{\sigma}_i^2 \quad (4)$$

The first term is the aforementioned epistemic uncertainty and it vanishes when $\hat{\theta}_i$ is a constant for all i , i.e. when we have zero parameter uncertainty. The second term represents the aleatoric uncertainty, namely the intrinsic noise in data.

3.4 ANCHORED ENSEMBLING

Anchored Ensembling (Pearce et al., 2020) represents the second candidate solution we adopt for applying BNNs in practice. The technique approximates Bayesian inference through an ensemble of NNs with parameters regularised around values drawn from an anchor distribution. This exploits a procedure which falls into a family of Bayesian methods called Randomised MAP Sampling (RMS). The RMS procedure follows the premise that addition of a regularisation term to a loss function returns a MAP parameter estimate, namely a point estimate of the Bayesian posterior. Injecting noise into this loss and sampling repeatedly (i.e. ensembling) results in a distribution of MAP solutions approximating the true posterior. Pearce et al. (2020) prove that under two conditions, partly present in NNs, their approximate method samples from the true posterior distribution. The authors show that their method is able to estimate an epistemic uncertainty extremely close to the one of a Gaussian Process (considered the ground truth). We particularly choose to employ this model for the expected superior capacity in providing well-calibrated uncertainty estimates. Furthermore, this model achieved competitive performance with other methods in supervised learning tasks and it has already been used for RL applications (Pearce et al., 2018).

3.5 GAUSSIAN PROCESSES

GPs (Williams & Rasmussen, 2006) are probabilistic non-parametric models widely used for applying Bayesian inference in regression tasks. GPs can be defined as an infinite set of random

variables, completely defined by a mean function μ and a covariance function (also called kernel) k : given $x, x' \in X$, $p(x) \sim \mathcal{GP}(\mu(x), k(x, x'))$. It is thus a distribution over functions. Function μ is often assumed equal to zero, since more complicated mean functions do not improve substantially the performance. When the GP has not yet observed the training data, the prior distribution lies around $\mu = 0$, according to the previous assumption. When the GP receives the information, the covariance matrix determines the functions, from the space of all possible functions, which are more probable. The conditioning property constrains this set of functions to pass through each training data. The GP model assigns a random variable to each point to predict and, through the marginalisation of each random variable, it computes the mean function value μ'_i and the standard deviation σ'_i of the i -th test point, where μ'_i is no longer equal to zero. If a test point is located on the training data, the function goes directly through it. Otherwise, the predictions are influenced by the training data proportionally to their distance, according to the kernel.

This model has been selected as it represents the principal counterpart of NNs for learning the dynamics of the environment in MBRL (Deisenroth & Rasmussen, 2011; Deisenroth et al., 2013). Since BNNs of infinite width are proved to converge to GPs (Neal, 2012), they are also considered the ground truth to evaluate the epistemic uncertainty estimation of approximate models. Moreover, its infinite capacity makes this model able to learn and memorize all observed data without the need of different training steps, in contrast to NN models (Gal et al., 2016). On the other hand, this model is known to scale cubically with number of data points and it is also negatively affected by problem dimensionality.

4 BENCHMARK PROCEDURE

This section describes the algorithm employed for the model comparison benchmark and it presents all implementation details of the considered models. A pure MBRL algorithm is usually composed of the general controller, the optimization algorithm, and the exploration policy. The choice of the MBRL algorithm is crucial for the success of the work. It should be tailored to serve our model comparison purpose, highlighting how well each model has learned the transition dynamics, while still delivering good performance. Considering this trade-off is key.

As a general controller, we employ the widely adopted *Model Predictive Control* (MPC) (Nagabandi et al., 2018; Chua et al., 2018; Wang et al., 2019), which is reported in Algorithm 1. The MPC controller is key to reach good performance even with an imperfect dynamics model. Its peculiarity is contained in Line 6, where only the first action from the optimal trajectory is executed. As a result, the agent re-plans the optimal sequence of actions at each time step, allowing to handle unexpected outcomes. Thus, the agent really interacts with the environment and a closed-loop policy is achieved.

Algorithm 1: Model Predictive Control

```

1 Run random policy  $\pi_0$  to collect data  $\mathcal{D} = \{(s_t, \mathbf{a}_t, s_{t+1})_i\}$ 
2 for  $K$  episodes do
3   Learn the dynamics model  $\hat{f}(s_t, a_t) = s_{t+1}$  to minimize some loss function.
4   for  $H$  steps do
5     Plan through  $\hat{f}(s, a)$  to choose actions.
6     Execute the first planned action only and observe the resulting state  $s_{t+1}$ 
7     Append  $(s_t, \mathbf{a}_t, s_{t+1})$  to  $\mathcal{D}$ 

```

4.1 LEARNING THE DYNAMICS MODEL

The dynamics model is learned in Line 3 of the MPC algorithm. How this is achieved depends on which model is employed among the classes introduced in Section 3. All NN-based models are composed by feed-forward fully-connected stacked layers and ReLU hidden activation functions. The Deterministic NN and the Deterministic Ensemble are trained with an MSE loss function:

$$\mathcal{L}_{\text{MSE}}(\hat{\theta}, \mathbf{x}, \mathbf{y}) = \frac{1}{N} \sum_{i=1}^N \|\mathbf{y}_i - \hat{\mathbf{y}}_i\|^2 \quad (5)$$

As in Kendall & Gal (2017), we assume a Gaussian likelihood and train the Concrete Dropout NN with the heteroscedastic loss:

$$\mathcal{L}_h(\hat{\theta}, \mathbf{x}, \mathbf{y}) = \frac{1}{N} \sum_{i=1}^N \frac{1}{\hat{\sigma}_i^2} \|\mathbf{y}_i - \hat{\mathbf{y}}_i\|^2 + \log \hat{\sigma}_i^2 \quad (6)$$

The Anchored Ensembling is instead trained with the regularised MSE loss function, as in Pearce et al. (2020):

$$\mathcal{L}_{\text{anc},j}(\hat{\theta}_j, \mathbf{x}, \mathbf{y}) = \frac{1}{N} \sum_{i=1}^N \|\mathbf{y}_i - \hat{\mathbf{y}}_{i,j}\|^2 + \frac{1}{N} \left\| \Gamma^{1/2}(\hat{\theta}_j - \theta_{\text{anc},j}) \right\|^2 \quad (7)$$

where the subscript j refers to the j -th NN of the ensemble, Γ is a diagonal regularisation matrix, and $\theta_{\text{anc},j}$ is the initialization drawn from the anchor distribution. For the GP model, the covariance function determines the learning performance. As suggested in Stein (1999) and Williams & Rasmussen (2006), we adopt the *Mat ern52* class.

Following (Deisenroth et al., 2013), for all models the training target \mathbf{y} is not directly the next state s_{t+1} , but the difference between the next state and the current state $\Delta_t = s_{t+1} - s_t$. The artifact was introduced to prevent GPs from falling to 0 in uncertain zones. This technique turned out to strongly reduce performance variance also in NNs and it has subsequently been adopted apart from GPs (Nagabandi et al., 2018). Lastly, all models are trained on a replay buffer (Lillicrap et al., 2020) which keeps only the last $N = 2000$ data and prevent the dataset \mathcal{D} from quickly containing an incredibly large amount of data. This trick also easily solves the issue of giving more importance to new trials than old ones (Gal et al., 2016). However, a well-known disadvantage is that the model might forget important old episodes.

In order to compute the predictive mean (Equation 2), we employ 5 NNs for the two ensembling methods. This is a number generally accepted to achieve satisfying results for both Deterministic (Beluch et al., 2018; Lakshminarayanan et al., 2016) and Anchored (Pearce et al., 2020) Ensembles. Concerning the Concrete Dropout NN, we follow the author’s implementation (Gal, 2018) and repeat the MC prediction $T = 20$ times.

4.2 THE OPTIMIZATION ALGORITHM

The core of the MBRL algorithm is composed by the online optimization algorithm that plans through the learned model (Line 5 of MPC). As in Nagabandi et al. (2018), we apply the Random Shooting (RS) method, which is reported in Algorithm 2.

Algorithm 2: Random Shooting

- 1 Set the number of trajectories N and the planning horizon n
 - 2 Pick $\mathbf{A}_1, \dots, \mathbf{A}_N$ from uniform distribution $p(\mathbf{A})$, where $\mathbf{A}_i = \{a_{0,i}, \dots, a_{n,i}\}$
 - 3 **for** trajectory $i = 1$ to N **do**
 - 4 **for** step $t = 0$ to $n - 1$ **do**
 - 5 Predict $\hat{f}(s_{t,i}, a_{t,i}) = s_{t+1,i}$
 - 6 Evaluate the trajectory $J(\mathbf{A}_i) = \sum_{t=0}^n r(s_{t,i}, a_{t,i})$
 - 7 Select $\mathbf{A}_{\text{best}} = \arg \max_i J(\mathbf{A}_i)$.
-

Based on the findings in Nagabandi et al. (2018) and Wang et al. (2019), we set the number of trajectories N to 500 and the planning horizon n to 20.

RS has been selected as optimization algorithm as the shooting methods family does not make strong assumptions on the transition dynamics and the reward function. Furthermore, RS has recently shown competitive results with more sophisticated algorithms such as Cross-Entropy Method (CEM) when combined with complex models (Wang et al., 2019). Finally, RS performance heavily depends on model accuracy. As a result, models that have learned better transition dynamics achieve better performance. This makes RS perfect for our model comparison objective and allows us to consider the sum of rewards per episode achieved during the RL training as a natural estimate of model performance.

As a common practice in MBRL, we also assume to access the reward function of the environment in order to compute Line 6 of RS. This allows to separate the dynamics model from the task to accomplish. Main benefit is obtaining a general model which can be applied to solve different tasks by only changing the reward function (Nagabandi et al., 2018).

4.3 THE EXPLORATION POLICY

The last piece of the benchmarking algorithm is formed by the exploration policy. This is needed to visit new states and potentially improve the dynamics model. Here, we try to fully exploit the diverse model capacities. Consequently, we implement two different policies depending on whether the employed model is able to estimate the epistemic uncertainty. For the Deterministic NN, a simple ε -greedy policy (Azizzadenesheli et al., 2018) is used. For all other models, we conceive a simple but effective exploration policy, which maximizes the Information Gain (IG) from visiting new states. This follows from the approximation of the IG with the *prediction gain* (Bellemare et al., 2016). Intuitively, a state is more informative if it causes a large change in model parameters. Let $\mathcal{U}(\cdot)$ be the function that outputs the epistemic uncertainty given a state and an action. The IG-based policy works as follows:

Algorithm 3: IG-based exploration policy

- 1 Randomly sample N actions a_1, \dots, a_N
 - 2 **for** action $i = 1$ to N **do**
 - 3 Evaluate the epistemic uncertainty $\mathcal{U}(s_t, a_i)$
 - 4 Select $a_{IG} = \arg \max_i \mathcal{U}(s_t, a_i)$
-

The full MBRL algorithm is reported in pseudo-code in Algorithm 4.

Algorithm 4: The RL algorithm

Input: Set number of episodes \mathbf{K} , episode length H , ε probability to explore, number of actions sampled for IG exploration, number of trajectories and planning horizon n for Random Shooting.

Output: Sum of rewards list $\mathbf{R} = \{R_0, \dots, R_{\mathbf{K}}\}$, with $R_i = \sum_{t=0}^H r_t$.

Data: $\mathcal{D} = \{\}$ replay buffer with capacity 2000 steps.

- 1 Initialize \mathcal{D} with a random controller for one episode.
 - 2 **for** episode = 1 to \mathbf{K} **do**
 - 3 Train dynamics model \hat{f} on \mathcal{D} .
 - 4 **for** step $t = 0$ to H **do**
 - 5 **if** $p \sim U(0, 1) < \varepsilon$ **then**
 - 6 Do exploration to find a_{expl} . // IG exploration or random exploration
 depending on the model
 - 7 $a_t := a_{expl}$
 - 8 **else**
 - 9 Do Random Shooting to select $\mathbf{A}_{best} = \{a_0, \dots, a_n\}$.
 - 10 $a_t := a_0$
 - 11 Execute a_t and collect reward r_t .
 - 12 Update dataset $\mathcal{D} \leftarrow \mathcal{D} \cup \{(s_t, a_t, s_{t+1})\}$
-

5 RESULTS

This section presents the benchmark results, i.e., the MBRL performance of the considered models for the selected environments. The latter are composed by one OpenAI Gym (Brockman et al., 2016) and five PyBullet Gym (Ellenberger, 2018–2019) environments, listed in Appendix A. In order to access the reward functions (see Section 4.2), these are implemented as reported in Appendix A.1. The employed model configurations are outlined in Appendix B for all tasks to ease reproducibility and transparency of the results.

Table 1: Benchmark results: models performance are reported with mean returns and standard error over all different tasks. Best results are highlighted in bold.

Model	Pendulum	InvertedPendulum	InvertedDoublePendulum
Deterministic NN	-424.68 ± 25.76	-279.52 ± 39.41	-66.19 ± 2.08
Deterministic Ensemble	-352.50 ± 22.41	- 81.48 ± 18.98	-48.37 ± 1.82
Concrete Dropout NN	- 315.40 ± 18.70	-126.17 ± 29.67	- 34.95 ± 1.67
Anchored Ensembling	-407.99 ± 26.62	-289.99 ± 40.06	-58.54 ± 2.03
Gaussian Process	-417.49 ± 24.33	-115.58 ± 34.45	-53.03 ± 1.98

Model	Reacher	HalfCheetah	Hopper
Deterministic NN	-35.22 ± 0.54	-80.08 ± 1.63	-750.75 ± 41.27
Deterministic Ensemble	-34.05 ± 0.48	-71.32 ± 1.46	-601.73 ± 38.52
Concrete Dropout NN	- 31.22 ± 0.35	- 27.09 ± 1.72	- 368.25 ± 37.50
Anchored Ensembling	-34.58 ± 0.51	-81.67 ± 1.49	-715.41 ± 38.70
Gaussian Process	-31.80 ± 0.36	-78.18 ± 1.50	-765.24 ± 43.83

The performance for each model is summarized in terms of average returns and standard error over 5 different random seeds. Table 1 depicts the benchmarks results.

The Concrete Dropout NN delivers the best performance in 5 out of 6 environments. Especially in the most complex tasks, such as HalfCheetah and Hopper, it dramatically outperforms all other models. The results further prove the high robustness provided by the dropout technique as a Bayesian approximation. In addition, this is the only model able to learn both the aleatoric and the epistemic uncertainty. This supports the intuition in Lakshminarayanan et al. (2016) that learning both the two types of uncertainty leads to improved performance.

The Deterministic Ensemble takes second place in most of the environments and even delivers the highest performance in the InvertedPendulum problem. With significantly higher results than the Deterministic NN, it proves that ensembling does improve performance in MBRL.

The Gaussian Process delivers high performance in Reacher, which is a relatively complex environment, and in easier tasks such as InvertedPendulum. However, as the complexity of the task increases, it is no longer possible to appreciate satisfying results. Indeed, it completely fails in HalfCheetah and Hopper. It also reports strong instability across different random seeds.

Unexpectedly, the Anchored Ensembling does not show substantially better performance when compared to the Deterministic NN. Along with the GP, it is affected by high instability across different random seeds.

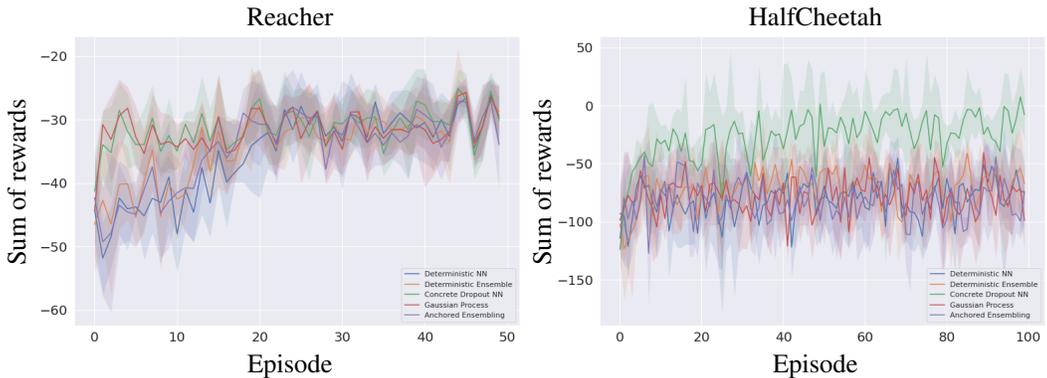


Figure 1: Model performance on Reacher (left) and HalfCheetah (right) environments.

Deepening our findings, we also show in Figure 1 two representative plots, whereby model performance is depicted by “Average Return” and “Standard Deviation of Returns”, as suggested in Islam et al. (2017), in Reacher and HalfCheetah environments. No smoothing technique is applied to the results shown. Learning curves denote the mean performance, while the shaded regions represent one standard deviation of rewards over 5 different random seeds. In the Reacher task, the Concrete Dropout NN and the Gaussian Process show an impressive sample-efficiency. While these two models converge to good results after only a couple of episodes, the others need about 20 episodes to achieve similar performance. In HalfCheetah, the Concrete Dropout NN clearly succeeds in terms of both sample-efficiency and overall performance, delivering unattainable results for all other models.

Lastly, it is noteworthy to consider the computing time of different models over the benchmarking comparison. As one might expect, the Deterministic NN is the fastest model across all tasks. Deterministic and Anchored Ensembles have a constant ratio with about 5 times longer computing time than the previous model. The Concrete Dropout NN and the Gaussian Process starts at a similar ratio than the former two in the easier tasks, but decay quickly as the complexity increases. In HalfCheetah, the computing times become about 15 and 30 times more than the Deterministic NN, respectively.

6 CONCLUSION

This work investigates the influence of models on the performance of RL algorithms. Model-based methods are expected to solve RL problems with fewer trials (i.e. higher sample-efficiency). However, the choice of model within the RL architecture has not formally been investigated in terms of its effect in improving the performance reported by novel algorithms. This work addresses this issue by comparing typical method classes that are commonly used as models for MBRL. The selected models are deterministic Neural Networks, ensembles of deterministic Neural Networks, Concrete Dropout Neural Networks, Anchored Ensembles, and Gaussian Processes. The model comparison is assessed on 6 benchmarking environments of increasing complexity.

Our results show that significant differences in model performance do exist. The Concrete Dropout NN reports persistently superior performance in 5 out of 6 benchmarking environments. Especially in the most complex tasks, it dramatically outperforms all other models. Its distinguishing strengths are higher sample-efficiency, robustness, and flexibility over different task complexities. On the negative side, this model requires larger computing time than other NN-based models. Without the need of different training steps to memorize all data, main strength of the GP is the remarkable sample-efficiency within the first handful of trials in low or medium dimensional problems. However, its lack of consistency prevents this model from converging to high performance compared to other explored models. As task complexity increases, it requires dramatically higher computing time and it is no longer possible to appreciate satisfying results. While the Deterministic Ensemble consistently achieves higher returns than the Deterministic NN, the differences are reduced in high dimensional problems. This suggests a lower flexibility in terms of task complexity when compared to the Concrete Dropout NN. The Anchored Ensembling is not generally able to outperform the Deterministic NN across the evaluated tasks. Lastly, the Deterministic NN does not deliver remarkable performance in any task but it succeeds in terms of computing time.

Our work proves that it is important that differences among models are taken into account. The improved performance reported by novel algorithms may (partly) be produced by the choice of a different modeling approach. Performance of novel MBRL algorithms, which introduce both a new RL approach and a different modeling choice, should be compared not only against previous methods, but also formally investigated with different modeling options in a similar fashion that we introduced in this work. Innovative techniques should therefore distinguish and make clear how much of the enhanced performance is due to improved algorithms or due to improved models.

The differences found also suggest that model choice can be tailored to the application of the RL algorithm. In the case where there are no computing time requirements, the Concrete Dropout NN is consistently the favorite candidate among those explored to deliver the highest sample-efficiency and overall performance, regardless of task complexity. However, if the task to accomplish is relatively low-dimensional and has to be solved only once with a few samples, also Gaussian Processes represent a valid alternative. On the other hand, if decisions must be made in the shortest possible time, a simpler Deterministic NN may be preferred.

REFERENCES

- Kai Arulkumaran, Marc Peter Deisenroth, Miles Brundage, and Anil Anthony Bharath. Deep reinforcement learning: A brief survey. *IEEE Signal Processing Magazine*, 34(6):26–38, 2017.
- Kamyar Aizzadenesheli, Emma Brunskill, and Animashree Anandkumar. Efficient exploration through bayesian deep q-networks. In *2018 Information Theory and Applications Workshop (ITA)*, pp. 1–9. IEEE, 2018.
- Marc Bellemare, Sriram Srinivasan, Georg Ostrovski, Tom Schaul, David Saxton, and Remi Munos. Unifying count-based exploration and intrinsic motivation. *Advances in neural information processing systems*, 29:1471–1479, 2016.
- William H Beluch, Tim Genewein, Andreas Nürnberger, and Jan M Köhler. The power of ensembles for active learning in image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9368–9377, 2018.
- Charles Blundell, Julien Cornebise, Koray Kavukcuoglu, and Daan Wierstra. Weight uncertainty in neural network. In *International Conference on Machine Learning*, pp. 1613–1622. PMLR, 2015.
- Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- Kurtland Chua, Roberto Calandra, Rowan McAllister, and Sergey Levine. Deep reinforcement learning in a handful of trials using probabilistic dynamics models. *arXiv preprint arXiv:1805.12114*, 2018.
- Ignasi Clavera, Jonas Rothfuss, John Schulman, Yasuhiro Fujita, Tamim Asfour, and Pieter Abbeel. Model-based reinforcement learning via meta-policy optimization. In *Conference on Robot Learning*, pp. 617–629. PMLR, 2018.
- Marc Deisenroth and Carl E Rasmussen. Pilco: A model-based and data-efficient approach to policy search. In *Proceedings of the 28th International Conference on machine learning (ICML-11)*, pp. 465–472. Citeseer, 2011.
- Marc Peter Deisenroth, Dieter Fox, and Carl Edward Rasmussen. Gaussian processes for data-efficient learning in robotics and control. *IEEE transactions on pattern analysis and machine intelligence*, 37(2):408–423, 2013.
- Benjamin Ellenberger. Pybullet gymperium. <https://github.com/benelot/pybullet-gym>, 2018–2019.
- Yarin Gal. Concretedropout, 2018. URL <https://github.com/yaringal/ConcreteDropout>.
- Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pp. 1050–1059. PMLR, 2016.
- Yarin Gal, Rowan McAllister, and Carl Edward Rasmussen. Improving pilco with bayesian neural network dynamics models. In *Data-Efficient Machine Learning workshop, ICML*, volume 4, pp. 25, 2016.
- Yarin Gal, Jiri Hron, and Alex Kendall. Concrete dropout. *arXiv preprint arXiv:1705.07832*, 2017.
- Geoffrey E. Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R. Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*, 2012.
- Riashat Islam, Peter Henderson, Maziar Gomrokchi, and Doina Precup. Reproducibility of benchmarked deep reinforcement learning tasks for continuous control. *arXiv preprint arXiv:1708.04133*, 2017.
- Michael Janner, Justin Fu, Marvin Zhang, and Sergey Levine. When to trust your model: Model-based policy optimization. *arXiv preprint arXiv:1906.08253*, 2019.

- Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? *arXiv preprint arXiv:1703.04977*, 2017.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. Simple and scalable predictive uncertainty estimation using deep ensembles. *arXiv preprint arXiv:1612.01474*, 2016.
- Stefan Lee, Senthil Purushwalkam, Michael Cogswell, David Crandall, and Dhruv Batra. Why m heads are better than one: Training a diverse ensemble of deep networks. *arXiv preprint arXiv:1511.06314*, 2015.
- Yuxi Li. Deep reinforcement learning: An overview. *arXiv preprint arXiv:1701.07274*, 2017.
- Timothy Paul Lillicrap, Jonathan James Hunt, Alexander Pritzel, Nicolas Manfred Otto Heess, Tom Erez, Yuval Tassa, David Silver, and Daniel Pieter Wierstra. Continuous control with deep reinforcement learning, September 15 2020. US Patent 10,776,692.
- David J. C. MacKay. A practical bayesian framework for backpropagation networks. *Neural Comput.*, 4(3):448–472, May 1992. ISSN 0899-7667. doi: 10.1162/neco.1992.4.3.448. URL <https://doi.org/10.1162/neco.1992.4.3.448>.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- Anusha Nagabandi, Gregory Kahn, Ronald S Fearing, and Sergey Levine. Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 7559–7566. IEEE, 2018.
- Anusha Nagabandi, Kurt Konolige, Sergey Levine, and Vikash Kumar. Deep dynamics models for learning dexterous manipulation. In *Conference on Robot Learning*, pp. 1101–1112. PMLR, 2020.
- Radford M Neal. *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media, 2012.
- Tim Pearce, Nicolas Anastassacos, Mohamed Zaki, and Andy Neely. Bayesian inference with anchored ensembles of neural networks, and application to exploration in reinforcement learning. *arXiv preprint arXiv:1805.11324*, 2018.
- Tim Pearce, Felix Leibfried, and Alexandra Brintrup. Uncertainty in neural networks: Approximately bayesian ensembling. In *International conference on artificial intelligence and statistics*, pp. 234–244. PMLR, 2020.
- Athanasios S Polydoros and Lazaros Nalpantidis. Survey of model-based reinforcement learning: Applications on robotics. *Journal of Intelligent & Robotic Systems*, 86(2):153–173, 2017.
- David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *nature*, 550(7676):354–359, 2017.
- ML Stein. Interpolation of spatial data. springer series in statistics. 1999.
- Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- Martijn Van Otterlo and Marco Wiering. Reinforcement learning and markov decision processes. In *Reinforcement Learning*, pp. 3–42. Springer, 2012.
- Tingwu Wang, Xuchan Bao, Ignasi Clavera, Jerrick Hoang, Yeming Wen, Eric Langlois, Shunshi Zhang, Guodong Zhang, Pieter Abbeel, and Jimmy Ba. Benchmarking model-based reinforcement learning. *arXiv preprint arXiv:1907.02057*, 2019.
- Christopher K Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.
- Zhi-Hua Zhou. *Ensemble methods: foundations and algorithms*. Chapman and Hall/CRC, 2019.

A ENVIRONMENTS

Table 2: Dimensionality, horizon, and number of episodes run of the environments.

Environment	Observation Space	Action Space	Horizon	Episodes
Pendulum	3	1	200	50
InvertedPendulum	4	1	100	35
InvertedDoublePendulum	11	1	100	100
ReacherPyBullet	9	2	50	50
Hopper	11	3	500	100
HalfCheetah	17	6	500	100

A.1 REWARD FUNCTIONS

Table 3: Reward functions of the environments.

Environment	Reward function
Pendulum	$-(((\theta_t + \pi) \% (2\pi)) - \pi)^2 - 0.1\dot{\theta}_t^2 - 0.001a_t^2$
InvertedPendulum	$-\theta_t^2$
InvertedDoublePendulum	$-\theta_t^2 - \gamma_t^2 - 0.001\dot{\theta}_t^2 - 0.005\gamma_t^2$
Reacher	$-distance_t - \ \mathbf{a}_t\ ^2$
Hopper	$\dot{x}_t - 0.1\ \mathbf{a}_t\ ^2 - 3h_t^2 + 1$
HalfCheetah	$\dot{x}_t - 0.1\ \mathbf{a}_t\ ^2$

B HIDDEN SIZE

Table 4: Hidden size of NN-based models.

Model	Pendulum	InvPendulum	InvDoublePendulum
Deterministic NN	2×32	2×40	2×100
Deterministic Ensemble	2×32	2×40	2×100
Concrete Dropout NN	2×100	2×100	2×200
Anchored Ensembling	2×40	2×40	2×100
Model	Reacher	HalfCheetah	Hopper
Deterministic NN	2×200	3×250	3×250
Deterministic Ensemble	2×200	3×250	3×250
Concrete Dropout NN	2×500	3×1024	3×1024
Anchored Ensembling	2×200	2×500	2×500