



UvA-DARE (Digital Academic Repository)

Uncertainty Aware Learning from Demonstrations in Multiple Contexts using Bayesian Neural Networks

Thakur, S.; van Hoof, H.; Gamboa Higuera, J.C.; Precup, D. ; Meger, D.

DOI

[10.1109/ICRA.2019.8794328](https://doi.org/10.1109/ICRA.2019.8794328)

Publication date

2019

Document Version

Final published version

Published in

2019 International Conference on Robotics and Automation (ICRA)

License

Article 25fa Dutch Copyright Act

[Link to publication](#)

Citation for published version (APA):

Thakur, S., van Hoof, H., Gamboa Higuera, J. C., Precup, D., & Meger, D. (2019). Uncertainty Aware Learning from Demonstrations in Multiple Contexts using Bayesian Neural Networks. In *2019 International Conference on Robotics and Automation (ICRA) : Montreal, Quebec, Canada, 20-24 May 2019* (Vol. 1, pp. 768-774). IEEE.
<https://doi.org/10.1109/ICRA.2019.8794328>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Uncertainty Aware Learning from Demonstrations in Multiple Contexts using Bayesian Neural Networks

Sanjay Thakur¹, Herke van Hoof², Juan Camilo Gamboa Higuera¹, Doina Precup¹, and David Meger¹

Abstract—Diversity of environments is a key challenge that causes learned robotic controllers to fail due to the discrepancies between the training and evaluation conditions. Training from demonstrations in various conditions can mitigate—but not completely prevent—such failures. Learned controllers such as neural networks typically do not have a notion of uncertainty that allows to diagnose an offset between training and testing conditions, and potentially intervene. In this work, we propose to use Bayesian Neural Networks, which have such a notion of uncertainty. We show that uncertainty can be leveraged to consistently detect situations in high-dimensional simulated and real robotic domains in which the performance of the learned controller would be sub-par. Also, we show that such an uncertainty based solution allows making an informed decision about when to invoke a fallback strategy. One fallback strategy is to request more data. We empirically show that providing data only when requested results in increased data-efficiency.

I. INTRODUCTION

This paper addresses the problem of learning from demonstration (LfD) when the environment changes frequently and outside the control of the robot. This situation is common in real-world deployments, such as self driving cars that must operate at all times of the day, in many traffic conditions and in all seasons. The performance of an LfD agent on an environment is dependent on the coverage and diversity of previous demonstrations. Training over multiple conditions has been shown to result in controllers that are more robust to such situations [1]. Furthermore, by making the policy dependent on the history, policies can rapidly adapt to different conditions [2], [3]. These robustness methods work up to a point, but cannot be expected to accommodate drastic and unexpected changes in the environment. As such, we think the agent should be able to signal conditions where it is not confident, allowing it to switch to a fall-back strategy.

We propose using Bayesian Neural Networks (BNNs) [4], [5] to learn robotic control policies from demonstration to model complex control functions while simultaneously assessing the uncertainty in their own predictions. We consider the case where task variables are not directly observed. For example, we may need to take several strides in the morning before determining that a joint is stiff today and will require additional force for optimal gait. Therefore, we train our networks on temporal windows of experience and observe that this allows effective assessment of uncertainty across different variations of a task which we will call *contexts*. In this setup, high predictive uncertainty indicates that the robot is currently encountering a context on which the

controller has not been sufficiently trained. A low predictive uncertainty, on the other hand, reflects high confidence that the learned policy will be successful in the current context. Our contributions in this work are

- to design a framework to generate predictive uncertainties over multiple task contexts as a confidence measure in an LfD setting scalably, and
- to propose an approach for actively requesting training data where the controller is not confident.

We support our claims through our experiments on a real and simulated robots. We first show that BNNs scale up easily to higher-dimensional tasks. We then show that there exists an inverse relation between the predictive uncertainty given by the BNN and the performance of the controller, so, the uncertainty can be used to predict its success. Finally, as fallback strategy we propose requesting more demonstrations when the predictive uncertainty is too high. We show that seeking more context specific demonstrations only when the context is highly uncertain leads to data-efficient learning.

In the next section, we define the objectives of our approach. We then describe related prior work such as active LfD and learning across multiple task variants. Next, we formalize our problem set-up and discuss relevant background. Subsequently, we outline how we use BNNs to represent uncertainty and how we leverage the uncertainty in the context of active LfD on multiple task contexts. We demonstrate the viability and benefits of our approach in section VII on a physical pendulum swing-up task, as well as several simulated robotics tasks.

II. OBJECTIVES AND NOTATION

We denote by D the set of task contexts that a robot might encounter one after another sequentially with each context d defined through a reward function $R_d(s_t, \mathbf{a}_t)$ and state transition dynamics $T_d(s_t, \mathbf{a}_t)$ which take as arguments an action \mathbf{a}_t and a state s_t . Each context d has a demonstrator π_d^* which provides demonstrations on d upon request, which are merged with the complete dataset \mathcal{D} . This demonstrator is (approximately) optimal with respect to the cumulative reward obtained on context d over the horizon H number of steps, $\sum_{t=1}^H R_d(s_t, \mathbf{a}_t)$.

As an example of a fallback strategy, we maintain a set-up where for each context, the controller which is being trained can halt execution at any point and ask for more demonstrations. The objective of the controller is to maximize the overall cumulative reward obtained on all the contexts in D while limiting the number of requests for demonstrations. In order to do so, the controller should have the ability to predict

¹School of Computer Science, McGill University

²Informatics Institute, University of Amsterdam

whether it will be able to succeed at the current context without additional demonstrations.

III. RELATED WORK

Our work aims to develop a method that allows a controller trained on demonstrations in several contexts to predict in which new contexts it might perform poorly. If this were the case, we propose that the learner should actively ask for more demonstrations under the current context. In this section, we will thus cover some related work on the topics of active learning from demonstration, approaches to learning multiple task variants, and active learning of control strategies based on uncertainty or confidence scores.

A. Active Learning From Demonstrations

LfD makes learning complex control applications [6]–[9] easy but demonstrations are expensive. Active learning aims to ameliorate this by seeking to select the most informative learning experiences [10]. We will show in subsection VII-B that our work leverages the predictive uncertainty as a cheap proxy as proposed in [11] to gauge what is informative. The work in [12], [13] and [14] quantify their uncertainty and detect unfamiliar and ambiguous states through an additional mechanism, e.g. depending on formulated pre- and postconditions [13]. Such an approach can be hard to scale up. Our proposed approach uses a single neural network that both controls the system and also detects unfamiliar states that makes it more scalable.

B. Learning Task Variants

Subtle changes between train-time and test-time conditions might cause controllers to fail silently and unpredictably [11] as highlighted in semi-supervised learning under domain shift [15], robust generalization [16], multi-task learning [17] and transfer learning [18]. The different conditions can thus be modelled as different (but closely related) *tasks*. We describe such different conditions by the term *contexts*.

Transfer learning addresses how policies can be transferred positively [18], [19] when there are subtle changes between train-time and test-time conditions which might cause controllers to fail silently [11], [15], [16], [18]. There have been attempts like allowing the simulator policy to initialize the real world policy [20], to allow policy adjustment model that compensates for simulator-real world discrepancies [21], training a RNN across multiple MDPs, to obtain useful biases that will aid the learning of other tasks in the future [2], [3]. However, these approaches lack an easily scalable way to detect training-testing distribution discrepancies and quantify confidence in a transfer.

C. Leveraging Uncertainty for Active Learning of Control

PILCO [22] uses Gaussian processes for learning a probabilistic dynamics model and long-term planning. However, Gaussian processes with generic kernels typically do not work well with high-dimensional inputs. In Deep-PILCO [23], [24] a dropout training approach [4] was used to quantify the uncertainty of a learned neural network

dynamics model. As an alternative, a heuristic measure has been used to quantify trust on the demonstrations in [25]. However, unlike in our work, all these methods focused on a single task, and most of them focus on model-uncertainty and not policy-uncertainty.

D. Learning sequential information

There are four broad classes of methods for capturing sequential information. The easiest to use with any off-the-shelf supervised learning methods, which we also employ in our work is using a *temporal windows* on sequences [26], [27]. The second method is using *recurrent models* such as LSTMs which are relatively difficult to train and work with evenly sampled sequences [28]. The third technique which works well with unevenly sampled data is *regression-based extrapolation* where training data is fitted to a curve [29], [30]. Lately, *Convolution based NN* techniques have been shown to perform at least as good as recurrent models for learning from sequences [31], [32], [33], [34], [35].

IV. BACKGROUND

The methods we describe in this paper require controllers π that provide a measure of uncertainty for their outputs, which can be used by an agent to assess its confidence about applying an action at a given state. Ideally, we would like a measure of uncertainty that is data-dependent; i.e. uncertainty should be higher for states \mathbf{s} where the agent has no data. This can be obtained by applying Bayesian inference. Given a dataset of demonstrations $\mathcal{D} = \{(\mathbf{s}_i, \pi_d^*(\mathbf{s}_i))\}$, we would like to estimate the posterior distribution over policies $p(\pi|\mathcal{D})$ to make predictions at possibly unseen states \mathbf{s}_* . This posterior distribution induces uncertainty on predictions as

$$p(\mathbf{a}_*|\mathbf{s}_*) = \int p(\mathbf{a}_*|\mathbf{s}_*, \pi)p(\pi|\mathcal{D})d\pi. \quad (1)$$

Two state-of-the-art methods that follow this approach are Gaussian Process (GP) regression and Bayesian Neural Networks (BNN). In this work, we will compare GP regression with a specific method for BNNs: Bayes-by-Backprop.

A. Gaussian Process Regression

Gaussian Processes are a popular *non-parametric* method for solving regression problems like the one in equation 1. GP regression proceeds by assuming that different evaluations of $\pi(\mathbf{s})$ for $\mathbf{s} \in \mathcal{D}$ are jointly Gaussian. This makes evaluating the integral in Eq. 1 tractable. GP regression requires specifying a *mean function* $m(\mathbf{s})$ and a *kernel function* $k(\mathbf{s}, \mathbf{s}')$. The mean functions acts as a prior for π , while the kernel function models the covariance of the outputs $\pi(\mathbf{s}), \pi(\mathbf{s}')$ given the inputs \mathbf{s}, \mathbf{s}' . In this case, the parameters of the policy θ correspond to the hyper-parameters of $m(\mathbf{s})$ and $k(\mathbf{s}, \mathbf{s}')$. Fitting is done by optimizing θ to maximize the marginal likelihood $p(\mathbf{a}|\mathbf{s})$ for all $\mathbf{s}, \mathbf{a} \in \mathcal{D}$.

Since the outputs are assumed to be jointly Gaussian, the resulting predictive distribution $p(\mathbf{a}_*|\mathbf{s}_*)$ is a Normal distribution $\mathcal{N}(\mathbf{a}_*|\boldsymbol{\mu}(\mathbf{s}_*), \boldsymbol{\sigma}^2(\mathbf{s}_*)\mathbf{I})$, where the parameters $\boldsymbol{\mu}(\mathbf{s}_*), \boldsymbol{\sigma}^2(\mathbf{s}_*)$ are calculated exactly by integrating (1) analytically. We use $\boldsymbol{\mu}(\mathbf{s}_*)$ as the action to be applied by the

agent and $\sigma^2(\mathbf{s}_*)$ as a measure of its uncertainty. Although inference can be done exactly, this method has a higher computational cost than inference with BNNs. We refer the reader to [36] for details on how to compute these quantities.

B. Bayes-by-Backprop

Bayes-by-Backprop [5, BBB] is an approximate variational inference method for training BNNs. The posterior distribution in this case is defined over the parameters of the model, $P(\pi|\mathcal{D}) = P(\mathbf{w}|\mathcal{D})$, where \mathbf{w} are the weights of the neural network model. Using the true posterior distribution for making predictions with neural networks is generally intractable. Approximate variational inference algorithms like BBB introduce an approximate posterior $q_\theta(\mathbf{w})$ which is easy to evaluate, and predictions are done via Monte Carlo sampling:

$$p(\mathbf{a}_*|\mathbf{s}_*) \approx \frac{1}{M} \sum_{i=1}^M p(\mathbf{a}_*|\mathbf{s}_*, \pi_{\mathbf{w}_i}), \quad \mathbf{w}_i \sim q_\theta(\mathbf{w}) \quad (2)$$

Fitting the model is done by minimizing the Kullback-Leibler divergence between the true and approximate posterior, which may be expressed with an objective of the form

$$\mathcal{L}(\mathcal{D}, \theta) = -\mathbb{E}_{q_\theta(\mathbf{w})} [p(\mathcal{D}|\mathbf{w})] + \text{KL}(q_\theta(\mathbf{w})|p(\mathbf{w})) \quad (3)$$

where $p(\mathcal{D}|\mathbf{w})$ is the likelihood of the data given the model, and $p(\mathbf{w})$ is a user-selected prior distribution over the weights. By making use of the reparameterization-trick [5], [37], the objective (3) can be approximated with Monte Carlo samples. This can be optimized via stochastic gradient descent with gradients evaluated by back-propagation. BBB sets $q_\theta(\mathbf{w})$ as a Gaussian distribution with diagonal covariance $q_\theta(\mathbf{w}) = \mathcal{N}(\mathbf{w}|\boldsymbol{\mu}, \boldsymbol{\sigma}) = \prod_{w_j \in \mathbf{w}} \mathcal{N}(w_j|\mu_j, \sigma_j)$. A suitable prior distribution in this case is a zero-mean Gaussian with diagonal unit covariance; $p(\mathbf{w}) = \mathcal{N}(\mathbf{0}, \mathbf{I})$.

Once the parameters θ^* that minimize (3) have been obtained, predictions are done via equation (2). This can be done by approximating $p(\mathbf{a}_*|\mathbf{s}_*)$ as a Normal distribution $\mathcal{N}(\mathbf{a}_*|\boldsymbol{\mu}(\mathbf{s}_*), \boldsymbol{\sigma}^2(\mathbf{s}_*)\mathbf{I})$, where the parameters $\boldsymbol{\mu}(\mathbf{s}_*)$, $\boldsymbol{\sigma}(\mathbf{s}_*)$ are estimated empirically from multiple evaluations of the model $\pi_{\mathbf{w}_i}(\mathbf{s}_*)$, with $\mathbf{w}_i \sim q_{\theta^*}(\mathbf{w})$. Again, the predictive mean $\boldsymbol{\mu}(\mathbf{s}_*)$ is used as the action applied by the agent, while the predictive variance $\boldsymbol{\sigma}^2(\mathbf{s}_*)$ provides a data-dependent measure of uncertainty.

V. LEARNING FROM DEMONSTRATIONS WITH BAYESIAN NEURAL NETWORKS ON MULTIPLE CONTEXTS

In this section, we first explain how we quantify uncertainty. We then build on this to perform active learning on multiple contexts in the LfD framework. In short, we use the predictive standard deviation of BBB to quantify the confidence in the controller performing well on the current context. High uncertainty means low confidence that following the policy will be successful.

To detect non-generalizable contexts, we use BBB to learn policies for which the inputs are temporal windows of the interaction history (see subsection V-A). The associated

uncertainty is used to decide whether the agent is confident enough that it can perform well in the current context, or if it should request demonstrations for additional training on the current context (subsection V-C).

The novelty in our mechanism lies in three parts that interact with each other, and together allow the application of Bayesian methods to quantify uncertainty in LfD policies: temporal-windows for input representation, learning policies for generating both actions for control and uncertainty over a given context, and the detector that decides whether the current policy generalizes well to the current context. These three elements are described in detail in the sections below, and summarized in Algorithm 1.

A. Temporal Windows

In order to implicitly identify differences in contexts and adapt to those changes, we use short histories of k time-steps or *temporal windows* instead of just a single state as input to the policy. For example, if the context dynamics depends on the mass of a system, the greater inertia can only be identified when comparing two subsequent states while a known force is applied.

The dataset \mathcal{D} for training is extracted from the demonstrations with such temporal windows, with a stride of one time-step. The overlapping windows \mathbf{x}_t consist of the last k states $\mathbf{s}_{t-k:t}$ paired with their associated actions $\mathbf{a}_{t-k:t-1}$ and rewards $r_{t-k:t-1}$.

B. Learning BNN Policies

The dataset is created as described in subsection V-A and the BBB is trained by minimizing Eq. 3. Hence, the parameters θ^* of the learned controller $\hat{\pi}$ are defined as

$$\theta^* = \arg \min_{\theta} \mathcal{L}(\mathcal{D}, \theta).$$

After optimization, the controller computes its prediction of the expert action and its uncertainty about the current situation with \mathbf{x}_t as input, as:

$$\mathbf{a}_t, \boldsymbol{\sigma}_t = \hat{\pi}_{\theta^*}(\mathbf{x}_t), \quad \text{with } \mathbf{a}_t = \mathbb{E}[\hat{\mathbf{y}}|\mathbf{x}_t], \quad \boldsymbol{\sigma}_t = \sqrt{\text{Var}[\hat{\mathbf{y}}|\mathbf{x}_t]}.$$

which are estimated using Monte-Carlo samples. The vector $\boldsymbol{\sigma}_t$ is then converted to a scalar σ_t by taking the average over the standard deviations obtained for each action-dimension.

C. Detector

The robot will encounter situations with contexts across the whole spectrum between being similar enough for good generalization to being quite different. Based on the controller's confidence, the agent decides whether to continue, or to stop and ask for demonstrations of the current context.

Particularly, we first smooth σ_t with a moving average over m time steps. The controller deems itself to be not confident enough to proceed if the average uncertainty exceeds a scaled adaptive threshold $c\Omega$ where c is a scaling factor and Ω is the adaptive threshold. The value of adaptive threshold Ω is automatically set to the average uncertainty obtained over all the contexts the controller has trained itself on so far. Both the constants c and m can be tuned to yield more or less

Algorithm 1 Select action and train if necessary

Require: $\mathbf{x}_t, \mathcal{D}, \theta, d, c$

```
1:  $\mathbf{a}_t, \sigma_t \leftarrow \hat{\pi}_\theta(\mathbf{x}_t)$ 
2: if smoothed( $\sigma_t$ ) >  $c\Omega$  and  $t > t_{start}$  then
3:    $\mathcal{D} \leftarrow \mathcal{D} \cup \text{GetDemonstrations}(d)$ 
4:    $\theta^* \leftarrow \arg \min_\theta \mathcal{L}(\mathcal{D}, \theta)$            #relearn
5:    $t \leftarrow 0$            # re-start episode in current context
6:    $\Omega = \text{AverageUncertainty}()$  # adapt query threshold
7:    $\mathbf{a}_t, \sigma_t = \hat{\pi}_{\theta^*}(\mathbf{x}_t)$            # re-select action
8: end if
9: return  $\mathbf{a}_t, \mathcal{D}, \theta^*$ 
```

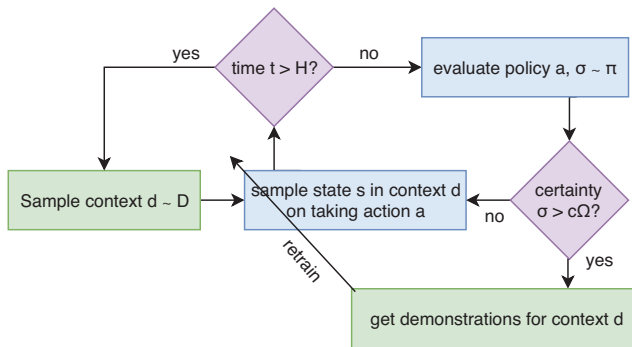


Fig. 1: The controller has to perform well on all tasks it faces sequentially with limited requests for task-specific demonstrations. A certainty threshold on the predictive variance is used to decide when to ask for demonstrations.

sensitive learners, which we demonstrate in subsection VII-C.2.

D. Connecting the dots

Algorithm 1 describes the steps of our approach, with Figure 1 showing how our method is used in an LfD setting. Using a BNN instead of a deterministic neural network allows us to obtain the predictive standard deviation along with the action for each temporal window. The temporal windows contain the robot’s state information over consecutive time-steps, which allows the BNN to implicitly infer the underlying dynamics. The uncertainty is passed to a detector that smooths this signal and compares it to a scaled threshold. When the uncertainty is higher than the scaled threshold, continuing will likely lead to poor performance. Hence, execution is paused and more training is done using new demonstrations solicited on the context at hand. The value of adaptive threshold Ω is initialized with 0 so that it always trains on requested demonstrations during the first context it faces. Following each training episode, the threshold Ω is adapted to the average uncertainty obtained on all the contexts that the controller has previously trained itself on.

VI. EXPERIMENTS

Our experiments¹ consist of training robust controllers for different task *contexts*. These closely related contexts

¹The code we used for these experiments is available at <https://git.io/fA1E4>

are obtained by varying variables of the environment, for example the masses and lengths of parts of the system to be controlled. These variables are not directly observable by the learning agent, but do have an effect on the outcome of applying an action. Since the learner is not told the context explicitly, it has to infer it from the past transitions within the temporal window. For all experiments, we require *expert* demonstrations. While our system can in principle operate with a human demonstrator, for purposes of reproducible science, we base our experiments on previously optimized agents. We use the definitions of episodic reward (r_d) and standard deviation (σ_d) on context d as $r_d = \sum_{t=0}^H r_t$ and $\sigma_d = \sum_{t=0}^H \sigma_t$ in our experiments.

A. Simulated environments

We tested our method on two modified OpenAI Gym environments [38]: HalfCheetah and Swimmer. In both environments, the goal is to find locomotion controllers to move forward as fast as possible. To simulate different contexts, we changed either or both the lengths and masses of various body parts like torso, middle, and back. For these experiments, we used proximal policy optimization (PPO) [39] to obtain expert policies for each context.

We also compare the relative learning performance of BBB to that of the GP on a simpler, low-dimensional task: a double integrator domain of a point mass sliding on a friction-less surface, where the goal is to apply forces so as to come to an halt exactly at coordinate $x = 0$. The demonstrator here is an LQR based controller obtained by solving the Riccati equation.

B. Physical robot experiment

We tested parts of our method on a physical Furuta pendulum [40]. The goal in this environment is to swing-up and balance the pendulum upright, by applying a torque to the base joint. This is analogous to the standard cart-pole swing-up task, with 4 continuous state dimensions and 1 continuous action dimension. We varied the pole mass to create different contexts. This was done by attaching weights of varying sizes to the pole. We used the model-based method described in [24] to obtain expert policies for each context.

VII. RESULTS

In this section, we will first show that our approach is more scalable and hence can do well even in high-dimensional tasks. We will then evaluate whether the confidence of our policy is a good proxy for success. Then, we will show how we can leverage the confidence to reduce the number of conditions in which demonstrations need to be given, and investigate the effect of the main tunable hyperparameters.

A. BNNs outperform GPs in high-dimensional tasks

We compare our BNN based proposed mechanism with GP in figure 2. Although GPs can do as well as BNNs in lower-dimensional domains, they fare significantly worse in higher-dimensions. In contrast, the number of input dimensions has an insignificant effect on the BNN’s ability to learn.

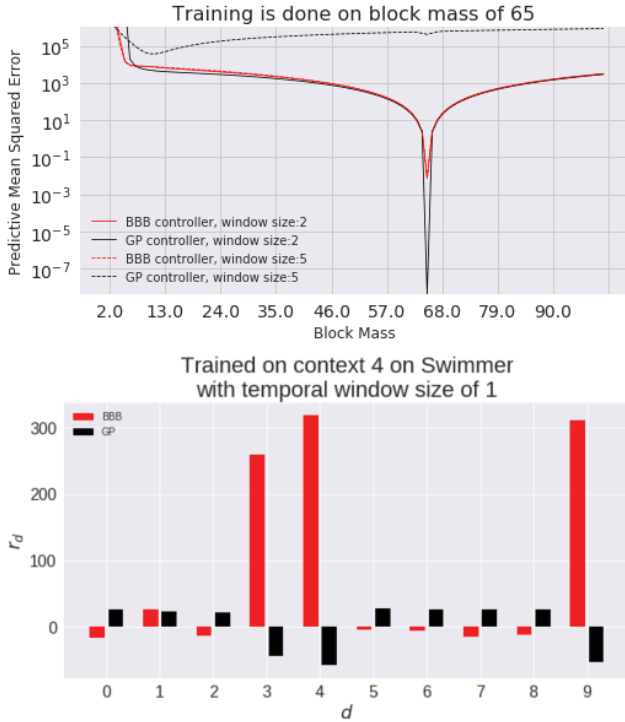



Fig. 2: GPs do not do as well as BBB when the number of input dimensions increases. Top figure shows how the performance of GP deteriorates with the increase in dimensions where the number of input dimensions corresponding to window sizes of 2 and 5 is 6 and 18 respectively. As a result, they never work with tasks that are already high dimensional (bottom figure) even with a temporal window size of 1, let alone use of deeper histories.

This is critical in most of the interesting real-world problems as they are high-dimensional.

B. Uncertainty as a predictor of success

In our work, the predictive uncertainty σ_d is used to detect and possibly preempt the execution on contexts the policy is unlikely to do well on. To test whether σ_d is indeed a good predictor for success, we compared the obtained σ_d and its corresponding empirical episodic reward r_d . We performed this experiment on the physical Furuta pendulum. Although the imitation learner does not attempt to optimize the reward, the demonstration policy does. Thus, high rewards are indicative of the learner closely following the demonstrator.

Figure 3b shows that σ_d is indeed related directly to the performance measurement r_d , so it can be used to predict whether the learner will do well. In particular, if the learner has been trained on contexts that are similar to the new context, it will tend to have a low predictive uncertainty and generalize well. If the new context is quite different, the policy will tend to have a high predictive uncertainty, and instead of attempting to perform the new context, the learner will first solicit new demonstrations for this context and then re-trains its policy on the full data-set.



(a)



(b)

Fig. 3: (a) Illustration of the experimental set-up for the robotic experimental platform. (b) For each swing-up task context, σ_d and r_d obtained during five independent runs are plotted against each other. Note the strong relationship between these quantities: higher σ_d is correlated with lower r_d . Thus, σ_d can be used as a predictor for task success. Training was performed on task context 6 in this case.

C. Data efficiency by deciding when to query

One way of validating if our proposed active learner invokes the fallback strategy effectively is by comparing the number of requests it makes for context-specific demonstrations with a *naive baseline* that solicits such context-specific demonstrations on every context it faces and a *random baseline* that never seeks any demonstration. We do two kinds of analysis for evaluating our active learner:

- 1) **Active LfD using predictive uncertainty sanity check:** Can the learner effectively ground its adaptive threshold in tandem with generating predictive uncertainty in any situation in order to correctly predict whether the controller can generalize to that context?
- 2) **Data-efficiency and tuning learner conservatism:** How judiciously the learner seeks demonstration and how does it affect the overall performance with respect to the *naive* and *random* baselines?

Answering these questions entails gathering data and evaluating policies over multiple task contexts. To keep experiments doable, we will do so in simulation rather than on the real robot system.

1) Active LfD using predictive uncertainty sanity check:

We first do a sanity check to see if our active learner has the ability to adapt its threshold to actively request context-specific demonstrations when it is most necessary. For this we purposefully arrange the sequence of contexts in a way

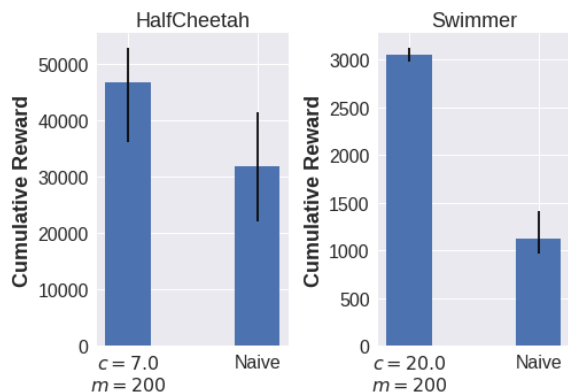


Fig. 4: Cumulative reward over all contexts in D after making 2 requests for demonstrations. Our active learner outperforms the naive learner that requests demonstrations in every context. Bars show average rewards over 3 episodes with error bars showing minimum and maximum values.

that the first two contexts are similar and different from the third. While the first two demonstration request for *naive* controller would be on the first and second context, our active learner’s first two requests should correspond to the first and the third context. Figure 4 shows that our proposed active learner achieves higher cumulative rewards over all the contexts in D than the *naive* controller after 2 requests for demonstrations which is a reflection of the fact that our active learner learns to perform well across a larger number of contexts than a naive controller with the same number of requests for demonstrations. The exact details of the ordering of the task can be found in the provided code.

2) Data-efficiency and tuning learner conservativeness:

After validating the potential of our active learner to make informed decisions about where to request for demonstrations, we now randomize the ordering of contexts in D to analyze data requirements for successfully solving all contexts in D . Figure 5 shows various configurations of our active learner using combinations of c and m to solve all D contexts. With a moderate number of requests for demonstrations, our proposed agents learn to solve the task adequately - appreciably better than *random* and close to the *naive* learner. hence attaining data-efficiency (minimizing the number of, possibly expensive, demonstrations). Note that the *naive* learner here solicits demonstrations on all the contexts it faces.

We would also like to note the effect of hyperparameters c (threshold scale) and m (amount of smoothing) on the conservativeness (asking for more demonstrations in case of doubt). Note that lower c and lower m leads to the agent asking for more demonstrations. These additional demonstrations offer greatly diminished returns, as the average reward attained did not go up appreciably as a result.

VIII. DISCUSSION AND FUTURE WORK

In this work, we proposed a mechanism based on Bayesian Neural Networks that can detect when task dynamics have changed significantly from the training dynamics. In a

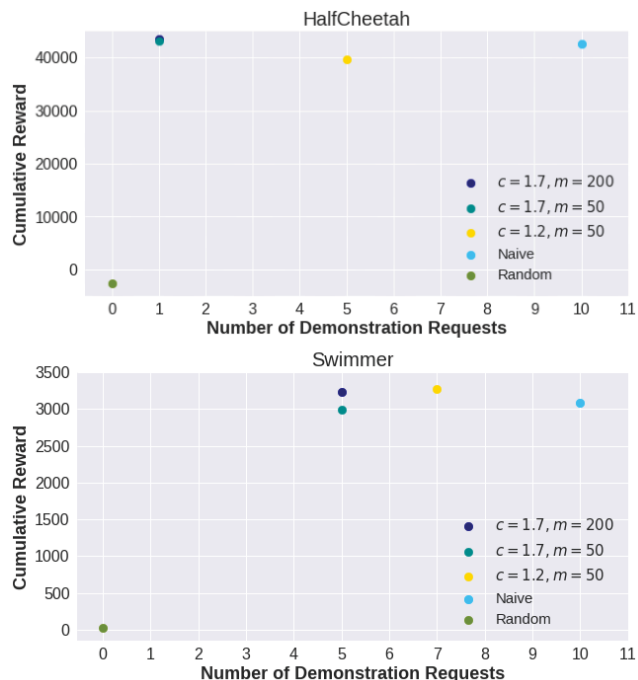


Fig. 5: The plots above show that lower c and m leads to seeking more context-specific demonstrations or more conservative behavior without much gain in reward. The plots also show that performance close to a *naive* learner that seeks demonstrations on every context, can be obtained by lesser number of such requests. The results were obtained by averaging over 5 simulation runs with different randomized orderings of contexts in D .

robotic pendulum swing-up task, we showed that the confidence of the learned controller obtained by our method is indeed a good predictor of success. Furthermore, as a proof of concept of effectively plugging in a fallback strategy, we showed how this mechanism can be used to judiciously request for demonstrations in novel conditions. The proposed Bayesian Neural Network scales better to a higher-dimensional task than a Gaussian process baseline.

We expect that our mechanism for detecting lack of confidence can be used to increase safety in robotics tasks. Depending on the task, our BNN controller could be combined with other fallback strategies such as a back-up controller or gracefully coming to a stop.

An avenue of future research is to combine our approach with dataset aggregation [41] to deal with the inherent LfD problem of covariate shift between the demonstrator’s and the learner’s state distribution. Furthermore, it is not always clear what aspects of a demonstration are critical for the desired behavior. We will investigate how to integrate a mechanism such as the one described by [42] to actively query a human supervisor and address this issue.

ACKNOWLEDGMENTS

We gratefully acknowledge the Natural Sciences and Engineering Research Council of Canada (NSERC) for funding and NVIDIA Corporation for the GPU donation.

REFERENCES

- [1] A. Rajeswaran, S. Ghotra, B. Ravindran, and S. Levine, "EPOpt: Learning robust neural network policies using model ensembles," in *ICLR*, 2017.
- [2] J. X. Wang, Z. Kurth-Nelson, D. Tirumala, H. Soyer, J. Z. Leibo, R. Munos, C. Blundell, D. Kumaran, and M. Botvinick, "Learning to reinforcement learn," CoRR, Tech. Rep. 1611.05763, 2016.
- [3] Y. Duan, J. Schulman, X. Chen, P. L. Bartlett, I. Sutskever, and P. Abbeel, "RL²: Fast reinforcement learning via slow reinforcement learning," CoRR, Tech. Rep. 1611.02779, 2016.
- [4] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in *ICML*, 2016, pp. 1050–1059.
- [5] C. Blundell, J. Cornebise, K. Kavukcuoglu, and D. Wierstra, "Weight Uncertainty in Neural Networks," in *ICML*, 2015.
- [6] B. D. Argall, S. Chernova, M. Veloso, and B. Browning, "A survey of robot learning from demonstration," *Robotics and autonomous systems*, vol. 57, no. 5, pp. 469–483, 2009.
- [7] P. Abbeel, A. Coates, and A. Y. Ng, "Autonomous helicopter aerobatics through apprenticeship learning," *Int. J. Robotics Research*, vol. 29, no. 13, pp. 1608–1639, 2010.
- [8] S. Calinon, *Robot programming by demonstration*. EPFL Press, 2009.
- [9] C. Finn, T. Yu, T. Zhang, P. Abbeel, and S. Levine, "One-shot visual imitation learning via meta-learning," CoRR, Tech. Rep. 1709.04905, 2017.
- [10] B. Settles, "Active learning literature survey," University of Wisconsin–Madison, Computer Sciences Technical Report 1648, 2009.
- [11] D. Amodei, C. Olah, J. Steinhardt, P. Christiano, J. Schulman, and D. Mané, "Concrete problems in AI safety," CoRR, Tech. Rep. 1606.06565, 2016.
- [12] D. H. Grollman and O. C. Jenkins, "Dogged learning for robots," in *ICRA*, 2007, pp. 2483–2488.
- [13] S. Tellex, R. A. Knepper, A. Li, T. M. Howard, D. Rus, and N. Roy, "Assembling furniture by asking for help from a human partner," in *Collaborative manipulation workshop at human–Robot interaction*, 2013.
- [14] S. Chernova and M. Veloso, "Interactive policy learning through confidence-based autonomy," *J. Artificial Intelligence Research*, vol. 34, no. 1, p. 1, 2009.
- [15] X. Zhu, "Semi-supervised learning literature survey," University of Wisconsin–Madison, Tech. Rep. 1530, 2005.
- [16] C. Zhang, S. Bengio, M. Hardt, B. Recht, and O. Vinyals, "Understanding deep learning requires rethinking generalization," in *ICLR*, 2016.
- [17] S. Ruder, "An overview of multi-task learning in deep neural networks," CoRR, Tech. Rep. 1706.05098, 2017.
- [18] S. J. Pan and Q. Yang, "A survey on transfer learning," *IEEE Trans. Knowledge Data Eng.*, vol. 22, no. 10, pp. 1345–1359, 2010.
- [19] A. Torralba and A. A. Efros, "Unbiased look at dataset bias," in *CVPR*. IEEE, 2011, pp. 1521–1528.
- [20] A. A. Rusu, M. Vecerik, T. Rothörl, N. Heess, R. Pascanu, and R. Hadsell, "Sim-to-real robot learning from pixels with progressive nets," CoRR, Tech. Rep. 1610.04286, 2016.
- [21] J. C. Gamboa Higuera, D. Meger, and G. Dudek, "Adapting learned robotics behaviours through policy adjustment," in *ICRA*, 2017, pp. 5837–5843.
- [22] M. Deisenroth and C. E. Rasmussen, "PILCO: A model-based and data-efficient approach to policy search," in *ICML*, 2011, pp. 465–472.
- [23] Y. Gal, R. McAllister, and C. E. Rasmussen, "Improving PILCO with Bayesian neural network dynamics models," in *Data-Efficient Machine Learning workshop, ICML*, 2016.
- [24] J. C. Gamboa Higuera, D. Meger, and G. Dudek, "Synthesizing neural network controllers with probabilistic model based reinforcement learning," *arXiv preprint arXiv:1803.02291*, 2018.
- [25] Z. Wang and M. E. Taylor, "Improving reinforcement learning with confidence-based demonstrations," in *IJCAI*, 2017.
- [26] R. J. Frank, N. Davey, and S. P. Hunt, "Time series prediction and neural networks," *Journal of intelligent and robotic systems*, vol. 31, no. 1-3, pp. 91–103, 2001.
- [27] M. S. Gashler and S. C. Ashmore, "Training deep fourier neural networks to fit time-series data," in *International Conference on Intelligent Computing*. Springer, 2014, pp. 48–55.
- [28] O. Nerrand, P. Roussel-Ragot, D. Urbani, L. Personnaz, and G. Dreyfus, "Training recurrent neural networks: Why and how? an illustration in dynamical process modeling," *IEEE Transactions on Neural Networks*, vol. 5, no. 2, pp. 178–184, 1994.
- [29] H. Drucker, C. J. Burges, L. Kaufman, A. J. Smola, and V. Vapnik, "Support vector regression machines," in *Advances in neural information processing systems*, 1997, pp. 155–161.
- [30] L. B. Godfrey and M. S. Gashler, "Neural decomposition of time-series data for effective generalization," *IEEE transactions on neural networks and learning systems*, vol. 29, no. 7, pp. 2973–2985, 2018.
- [31] S. Bai, J. Z. Kolter, and V. Koltun, "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling," *arXiv preprint arXiv:1803.01271*, 2018.
- [32] A. Van Den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, "Wavenet: A generative model for raw audio," in *SSW*, 2016, p. 125.
- [33] N. Kalchbrenner, E. Grefenstette, and P. Blunsom, "A convolutional neural network for modelling sentences," *arXiv preprint arXiv:1404.2188*, 2014.
- [34] Y. N. Dauphin, A. Fan, M. Auli, and D. Grangier, "Language modeling with gated convolutional networks," *arXiv preprint arXiv:1612.08083*, 2016.
- [35] J. Gehring, M. Auli, D. Grangier, and Y. N. Dauphin, "A convolutional encoder model for neural machine translation," *arXiv preprint arXiv:1611.02344*, 2016.
- [36] C. E. Rasmussen, "Gaussian processes in machine learning," in *Advanced lectures on machine learning*. Springer, 2004, pp. 63–71.
- [37] D. P. Kingma, T. Salimans, and M. Welling, "Variational dropout and the local reparameterization trick," in *Advances in Neural Information Processing Systems*, 2015.
- [38] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," CoRR, Tech. Rep. 1606.01540, 2016.
- [39] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," CoRR, Tech. Rep. 1707.06347, 2017.
- [40] K. Furuta, M. Yamakita, and S. Kobayashi, "Swing-up control of inverted pendulum using pseudo-state feedback," *Journal of Systems and Control Engineering*, vol. 206, pp. 263 – 269, 1992.
- [41] S. Ross, G. J. Gordon, and J. A. Bagnell, "No-regret reductions for imitation learning and structured prediction," in *AISTATS*. Citeseer, 2011.
- [42] D. Sadigh, A. D. Dragan, S. Sastry, and S. A. Seshia, "Active preference-based learning of reward functions," in *R:SS*, 2017.