

An Adversarial Objective for Scalable Exploration

Bernadette Bucher^{1*}, Karl Schmeckpeper^{1*}, Nikolai Matni¹, Kostas Daniilidis¹

Abstract—Collecting new experience is costly in many robotic tasks, so determining how to efficiently explore in a new environment to learn as much as possible in as few trials as possible is an important problem for robotics. In this paper, we propose a method for exploring for the purpose of learning a dynamics model. Our key idea is to minimize a score given by a discriminator network as an objective for a planner which chooses actions. This discriminator is optimized jointly with a prediction model and enables our active learning approach to sample sequences of observations and actions which result in predictions considered the least realistic by the discriminator. Comparable existing exploration methods cannot operate in many prediction-planning pipelines used in robotic learning without hardware modifications to standard robotics platforms in order to accommodate their large compute requirements, so the primary contribution of our adversarial exploration method is scalability. We demonstrate progressively increased performance of our adversarial exploration approach compared to leading model-based exploration strategies as compute is restricted in simulated environments. We further demonstrate the ability of our adversarial method to scale to a robotic manipulation prediction-planning pipeline where we improve sample efficiency and prediction performance for a domain transfer problem.

I. INTRODUCTION

Robotic learning solutions can accomplish complex tasks but struggle from expensive data collection requirements [1]–[10]. These model-based reinforcement learning methods are typically trained with randomly collected [1], [2], [5], [7], [11] or on-policy [6], [12], [13] data. Targeted exploration policies collecting new experiences for training allow these models to be trained on off-policy data with lower sample complexity than random collection [14], [15].

Methods for *curiosity* maximize expected information gain of predictive models which can be used to perform targeted sampling [16]–[18]. Model-free curiosity derives rewards *after* an action is taken and thus requires observation of the action outcome. This approach necessitates integration with model-free reinforcement learning in which rewards provide feedback to an updated policy for selecting actions. In contrast, model-based methods are active learning strategies which use a prediction model directly to select actions, so curiosity measurements can be made *before* the action is taken to execute curious behavior. The fundamental difference between these approaches is visualized in Figure 2.

The active learning and active perception literature has long established the ability of good sampling strategies to increase sample efficiency and model performance in robotics [19]–[21]. Model-based curiosity research distinguishes itself

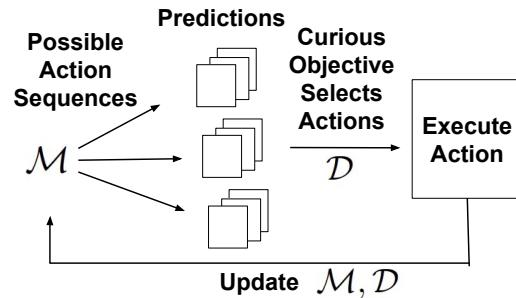


Fig. 1: Our adversarial objective for curiosity operates on potential trajectories generated by a predictive model. These predictions are evaluated with a discriminator and the trajectory that corresponds to the least realistic prediction is executed. The predictive model and the discriminator are updated with the newly collected data.

within this broader set of work by focusing on objectives for predictive model information gain. Existing model-based curiosity methods which operate on high dimensional and continuous state spaces rely on ensembles of prediction models to estimate these objectives. We propose the first model-based curiosity method which does not rely on the use of an ensemble.

In work most similar to our own, model-based active exploration (MAX) computes an estimate of the uncertainty of model predictions derived from the variance of outcomes computed within an ensemble of prediction models [14]. The compute requirements of this ensemble-based approach scale linearly with the number of models in the ensemble which can become intractable for compute intensive prediction models. For example, many modern vision-based prediction methods require the full capacity of a GPU to train which would require ensemble based approaches to use many GPUs to explore. This requirement is unreasonable for hardware constrained systems. In this work, we present an objective for scalable exploration based on minimizing a score given by a discriminator network in order to choose actions which result in outcomes considered the least realistic by our adversarial network visualized in Figure 1.

Our Contributions We present an adversarial curiosity reward which we use as an objective in model-based reinforcement learning systems. We integrate this method into two distinct prediction-planning pipelines. In the first pipeline, we perform state-based prediction and plan with a Markov Decision Process to generate exploration and exploitation policies. We demonstrate the scalability of our approach by comparing performance with MAX across varying compute

* Alphabetical ordering; the first two authors contributed equally.

¹GRASP Laboratory, University of Pennsylvania {bucherb, karls, nmatni, kostas}@seas.upenn.edu

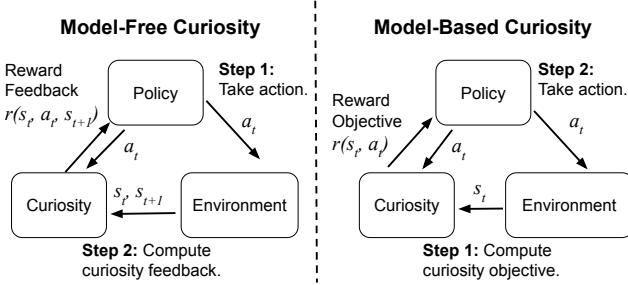


Fig. 2: Our model-based curiosity approach compared to established model-free formulations of curiosity. Existing model-free approaches to curiosity calculate the curiosity score after the action has been taken, while our approach uses the curiosity score to decide which action to take.

restrictions in simulation and show progressively increased performance of our method compared to MAX at lower compute targets. We also integrate our discriminator with a video prediction model and the cross-entropy method for planning, a compute intensive prediction-planning pipeline for robotic manipulation. This pipeline is trained with randomly sampled data in all prior work and would require hardware modifications on standard robotics platforms to be able to use MAX as an exploration policy. Our method improved sample efficiency without sacrificing task performance. In addition, our method increased prediction performance. To the best of our knowledge, this application is the first use of model-based curiosity for vision-based robotic manipulation.

Our paper is organized as follows. First, we review the literature on sampling with both curiosity and active learning methods. Then, we present our adversarial curiosity objective for sampling in a model-based reinforcement learning problem. Next, we favorably compare the ability of our method to complete downstream tasks under computation constraints with other curiosity methods using policies learned via exploration. Finally, we demonstrate the ability of our method to scale to a robotic manipulation pipeline in which we show increased sample efficiency and improved prediction performance in a domain transfer problem.

II. RELATED WORK

We specify our method as an adversarial curiosity objective which we use to perform active learning in model-based reinforcement learning. Thus, we contextualize our work in the active learning and curiosity literature. The goal of both active learning and curiosity is to select samples with which to train a model such that the model gains the most possible information from the sampled data.

Model-Free Curiosity Several existing methods use exploration incentives to estimate and seek visual novelty in model-free reinforcement learning. Approaches include rewarding policies based on state visit counts [22]–[24], observed error in a prediction model [25]–[28], disagreement in an ensemble [15], [29], and information gain in a Bayesian neural network [30]. We present a model-based curiosity method using the scores from a discriminator network. These

same scores could be used as feedback to a model-free method, but we only evaluate the model-based case in this work to highlight the key characteristic of our curiosity approach: scalability.

Model-Based Curiosity There are far fewer model-based curiosity methods than model-free. Discrete count-based methods [31] estimate the learning progress of a potential sample [32]. Ensemble-based approaches estimate the uncertainty [15], [33] or entropy [14] of predictions. Bayesian models are also used to minimize uncertainty [34]. Our work differs from previous model-based active exploration methods in that it is computationally efficient in high-dimensional continuous domains through the use of a discriminator which provides scores for the realism of our model predictions.

Adversarial Curiosity The concept of adversarial curiosity was first proposed prior to the integration of curiosity with planning algorithms [35]. This work suggested that the formulation of the minimax problem presented a method of encoding introspective behavior in a model. Minimax optimization problems such as the one we propose can be argued to provide intrinsic motivation for a model to invent novel information about which to learn [36]. This behavioral paradigm is considered a form of curiosity [16], [18], [35], [36]. We propose the first explicit model of adversarial curiosity with experimental evaluation in this work.

Active Learning Active learning is the process where a machine learning algorithm selects its training data to improve its data efficiency and performance [37]. Many methods for active learning have been proposed, including work most similar to our own which samples the most uncertain data points [38]–[40] and work most similar to our experimental baselines which sample where an ensemble of models disagrees [41]–[43]. Other active learning approaches propose sampling data that will cause the largest expected information gain [44], the largest expected change in the model [45]–[47], the largest expected reduction in variance [20], or the largest estimated error reduction [19], [48], [49]. To distinguish active learning from curiosity, active learning is specifically formulated to determine samples to select *before* sampling whereas the curiosity literature includes both active methods [34] and methods in which rewards can only be computed *after* sampling [15], [25]. Thus, the method we propose in this work is both a method for curiosity and active learning.

Summary The closest exploration rewards to which we compare our method in Section IV are prediction error [25] and entropy over ensemble predictions [14]. Following the notation in Figure 2, prediction error is computed as

$$r(s_t, a_t, s_{t+1}) = \|s_{t+1} - \hat{s}_{t+1}\|_2 \quad (1)$$

where \hat{s}_{t+1} is the estimated future state by a prediction model computed from s_t and a_t . However, computing a prediction error reward requires access to the future state s_{t+1} , so we can only use this reward with model-free reinforcement learning methods. Entropy over ensemble predictions is

computed via the Jensen-Renyi Divergence (JRD) as

$$r(s_t, a_t) = JRD\{\hat{s}_{0,t+1}, \dots \hat{s}_{N,t+1}\} \quad (2)$$

where $\hat{s}_{i,t+1}$ for $i \in [0, N]$ is the future state predictions from each of the $N + 1$ models in an ensemble computed from s_t and a_t . Therefor, this method is compatible with model-based reinforcement learning, but the performance of the method is dependent on the size of the ensemble. Our discriminator computes a reward

$$r(s_t, a_t) = \mathcal{D}(s_t, a_t, \hat{s}_{t+1}) \quad (3)$$

where \hat{s}_{t+1} is the estimated future state by a prediction model computed from s_t and a_t . The true future state s_{t+1} is used to train \mathcal{D} so that \mathcal{D} can differentiate between $(s_t, a_t, \hat{s}_{t+1})$ and (s_t, a_t, s_{t+1}) , but when executing exploration, only the predicted future state \hat{s}_{t+1} is needed to evaluate the reward, allowing it to be used with model-based approaches without requiring an ensemble of prediction models. If the discriminator can easily determine that a state \hat{s}_{t+1} is a model prediction versus a ground truth observation, this unrealistic trajectory generated by our model is given a low score by the discriminator. Thus, our exploration policy minimizes the reward in equation 3. The full methodological details of our approach are presented in Section III.

III. ADVERSARIAL CURIOSITY OBJECTIVE

Consider the dynamics of a system \mathcal{F} mapping a context window of C past states $s_{t-C:t}$, and H potential future actions $a_{t:t+H}$, to future states via

$$s_{t+1:t+H+1} = \mathcal{F}(s_{t-C:t}, a_{t:t+H}), \quad (4)$$

We then denote by

$$\mathbf{x} = (s_{t-C:t}, a_{t:t+H}, s_{t+1:t+H+1}) \quad (5)$$

any trajectory generated by system (4), and denote by $p(\mathbf{x})$ the distribution over these trajectories. To lighten notational burden going forward, we let $\mathbf{a} := a_{t:t+H}$, $\mathbf{c} := s_{t-C:t}$, $\mathbf{h} := s_{t+1:t+H+1}$.

Our model-based curiosity method is defined in terms of the following three components:

- 1) We approximate the system (4) with a model \mathcal{M} that generates predictions of future states \hat{s} given a context window of past states \mathbf{c} and a sequence of potential future actions \mathbf{a} . Thus, our prediction model is given by

$$\hat{s}_{t+1:t+H+1} = \mathcal{M}(\mathbf{c}, \mathbf{a}). \quad (6)$$

- 2) A discriminator \mathcal{D} , which assigns a score d_t to each real trajectory \mathbf{x} generated by system (4) as well as imagined trajectories generated by the prediction model (6). To train the discriminator \mathcal{D} jointly with a model \mathcal{M} , we solve the minimax optimization problem

$$\begin{aligned} \min_{\mathcal{M}} \max_{\mathcal{D}} \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [L(\mathbf{h}, \mathcal{M}(\mathbf{c}, \mathbf{a}))] + \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} [\log \mathcal{D}(\mathbf{x})] \\ + \mathbb{E}_{(\mathbf{c}, \mathbf{a}) \sim p(\mathbf{x})} [\log (1 - \mathcal{D}(\mathbf{c}, \mathbf{a}, \mathcal{M}(\mathbf{c}, \mathbf{a})))]. \end{aligned} \quad (7)$$

The first term in the objective function of optimization problem (7) computes the loss L for training the model \mathcal{M} , using whatever loss function would be used to train the model without curiosity. The second term captures the ability of the discriminator to identify realistic trajectories generated by system (4), whereas the third term simultaneously reflects the predictive ability of the model \mathcal{M} , as well as the ability of the discriminator \mathcal{D} to distinguish between real and imagined trajectories.

The inner maximization trains the discriminator \mathcal{D} to differentiate between trajectories sampled from the data distribution \mathbf{x} and predicted trajectories $(\mathbf{c}, \mathbf{a}, \mathcal{M}(\mathbf{c}, \mathbf{a}))$. The outer minimization optimizes the performance of the prediction model \mathcal{M} . In summary, this minimax problem sets up a competition in which the prediction model tries to learn to make good enough predictions to fool the discriminator while the discriminator tries to improve differentiation of predictions from data samples.

After \mathcal{D} is trained, the discriminator scores for our imagined trajectories are evaluated as

$$d_t = \mathcal{D}(\mathbf{c}, \mathbf{a}, \mathcal{M}(\mathbf{c}, \mathbf{a})). \quad (8)$$

- 3) With these pieces in place, we can now define the curiosity based optimization problem that we solve in order to select action sequences which optimize a curiosity objective defined in terms of the discriminator score. In particular, we define a planner \mathcal{P} that selects an action sequence \mathbf{a} of the H next actions which minimize the discriminator score by solving the optimization problem:

$$\mathcal{P}(\mathbf{c}, \mathbf{a}, \mathcal{M}, \mathcal{D}) := \arg \min_{\mathbf{a}} \mathcal{D}(\mathbf{c}, \mathbf{a}, \mathcal{M}(\mathbf{c}, \mathbf{a})) \quad (9)$$

It then follows that the actions resulting in the least realistic predictions are selected by the planner defined by optimization problem (9), resulting in qualitatively more *curious* behavior.

IV. EXPLORATION RESULTS UNDER COMPUTATION CONSTRAINTS

We compare our curiosity method to two model-based exploration methods [14].¹ Model-based active exploration (MAX) rewards exploration with the Jensen-Shannon or Jensen-Renyi divergence between an ensemble of prediction models to measure prediction uncertainty. The Jensen-Renyi divergence is used for continuous state spaces as is the case in our experiments. Trajectory variance active exploration (TVAX) computes the variance in sampled trajectories between an ensemble of prediction models as an exploration reward. We also compare our method to PERX, a model-free prediction error approach [25], and JDRX, a model-free version of MAX, as well as to random exploration.

¹For additional implementation details and to reproduce our results, please refer to our released code: <https://github.com/bucherb/adversarial-curiosity>

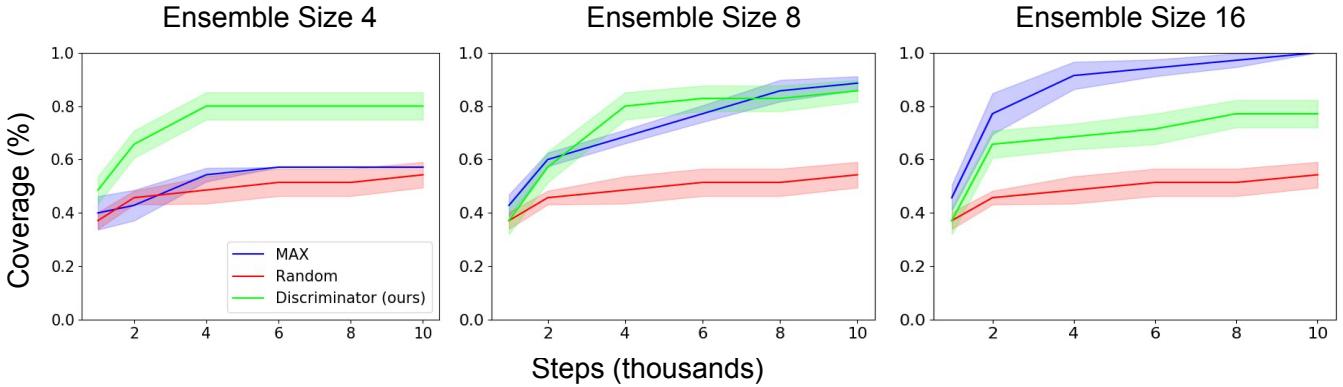


Fig. 3: Maze coverage results in the Ant Maze OpenAI environment averaged over 5 trials initiated with random seeds. Our approach outperforms ensemble-based methods when low ensemble sizes are used due to computation constraints. Perception models in robot learning applications often use the full capacity of a GPU to train, so the size of the ensemble would correspond with the number of needed GPUs making ensemble-based exploration strategies computationally intractable.

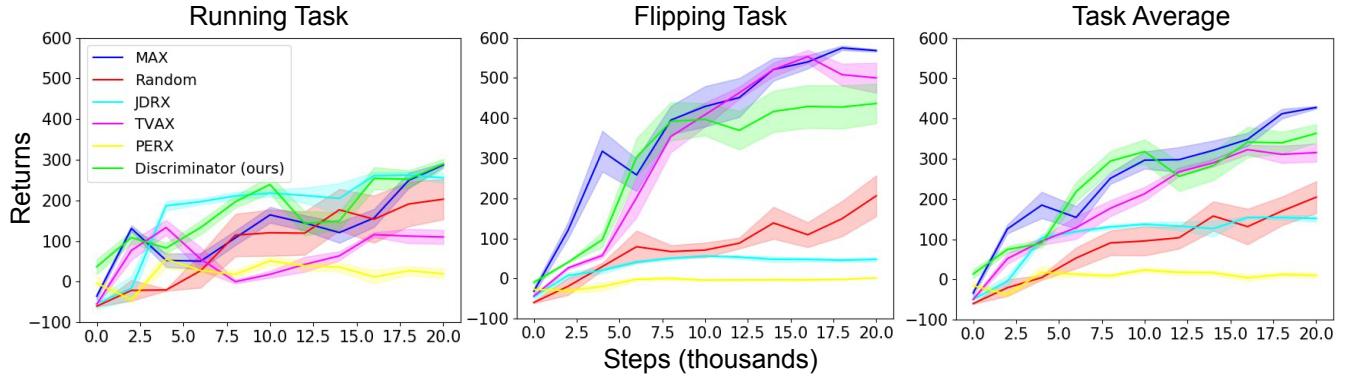


Fig. 4: Extrinsic rewards from running and flipping task completion in the Half Cheetah OpenAI environment averaged over 5 trials initiated with random seeds broken down by each distinct method for exploration used in training. All ensemble based curiosity methods (MAX, JDRX, TVAX) were trained with an ensemble size of 32. Our overall task performance is a close second to MAX even in this case where compute is not restricted.

In order to execute a fair comparison, we use the same prediction-planning pipeline proposed with the MAX algorithm across all of our simulated experiments so that performance differences can be based solely on the different exploration policy approaches. All methods use an ensemble of fully connected neural networks trained with a negative log-likelihood loss jointly with our discriminator to predict next state distributions. Next state predictions are given by the distribution mean. We emphasize that this use of ensembles as a predictive model is not necessary in the case of our curiosity method other than to standardize the prediction approach for comparison purposes whereas MAX, JDRX, and TVAX all require an ensemble-based prediction approach to derive their curiosity rewards. Similarly, the discriminator is only used for our curiosity method but trained with every method to standardize the prediction approach. The model-free methods use soft-actor critic (SAC) [50] to learn a policy. The model-based methods are trained with data gathered by an MDP using an exploration objective following the framework presented used originally by MAX

[14].

We first compare maze coverage results in the OpenAI Ant Maze environment in our experiments presented in Figure 3. We note that the key contribution of our approach comes from the scalability of the discriminator. We observe that at high ensemble sizes (16 or more prediction models), MAX outperforms our discriminator. However, as ensemble size decreases, our method significantly outperforms MAX. Thus, we show that when compute restrictions prevent the use of large ensembles of prediction models, the performance of ensemble based methods are significantly compromised while our method continues to perform well. We demonstrate the significance of this result in Section V where we integrate our method in an established prediction-planning pipeline for robotic manipulation for which hardware changes would be required to run an ensemble based approach.

Next, we evaluate the performance of our method when compute is not restricted. We learn running and flipping tasks in the OpenAI Half Cheetah environment which has an 18-dimensional continuous observation space and

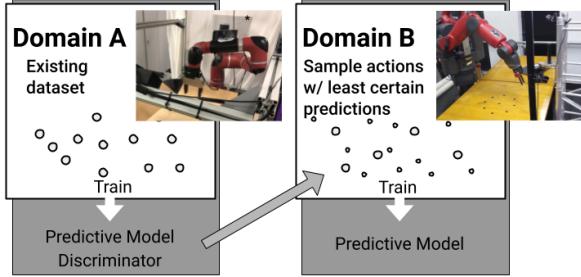


Fig. 5: Active sampling to enable domain transfer. In our robotics experiments, our method trains an action-conditioned prediction model and a discriminator on the dataset in the initial domain. It then samples actions from the new domain that result in the most uncertain predictions, allowing it to train a prediction model in the new domain with a small number of samples.

a 6-dimensional continuous action space [51]. Following MAX [14], we add Gaussian noise of $\mathcal{N}(0, 0.02)$ to the actions to increase the difficulty. Policies are trained purely with each exploration strategy and then used to execute downstream tasks. We compare the performance of each exploration method by evaluating achieved extrinsic rewards in these tasks over 5 randomly seeded trials and present them in Figure 4. We find that our method either slightly underperforms or performs comparably to all the baselines in both downstream tasks: running and flipping. However, we note that these comparisons were evaluated using an ensemble of 32 models, which is the optimal number for MAX [14].

V. ROBOTIC MANIPULATION

We demonstrate the ability of our adversarial curiosity method to meaningfully scale to robotics problems by integrating our approach in a prediction-planning pipeline commonly studied robot learning approaches to manipulation. Large amounts of data are required for training robot learning models which is time consuming to collect. To address this challenge, recent research demonstrated the ability of a class of these data-driven methods to transfer across platforms [7]. This capability allows for publicly available datasets to be leveraged for the bulk of the training, and only a small amount of data needs to be collected for fine-tuning in the new domain. We demonstrate increased sample efficiency and gains in prediction performance from incorporating our discriminator in this pipeline to perform targeted sampling as illustrated in Figure 5. This pipeline modification and improvement was not possible with existing model-based curiosity methods including MAX without hardware modifications. All prior methods rely on the use of ensembles, and pipelines leveraging video prediction models typically utilize the full capacity of the GPU as is the case for the pipeline we use in our experiment.

States are RGB images in this set of experiments. We use a variant [7] of the Visual Foresight video prediction model [5]. A stack of convolutional LSTMs is used to predict

a flow field from an image s_t and action a_t . This flow field is then applied directly to the input image s_t to predict the next image frame \hat{s}_{t+1} . The true next image frame s_{t+1} is observed after the given action is taken. This network is optimized with an L_1 loss between the predicted image \hat{s}_{t+1} and true image s_{t+1} . In practice, these models perform predictions out to some horizon H using a context of C image frames in which the flow field estimates are applied recursively across the prediction horizon.

In the training procedure described by Step 1 in Figure 6, our prediction model is optimized jointly with the discriminator defined in Section III. This is similar to the loss used by Stochastic Adversarial Video Prediction [52], where the combination of prediction error and an adversarial loss were shown to improve prediction quality and convergence. We use this prediction model and the cross-entropy method (CEM) [53] to optimize for a sequence of actions that minimize Equation 9. This procedure causes the robot to select a sequence of actions that generate unrealistic looking predictions.

A. Sampling Analysis

We evaluate the ability of our curiosity objective to effectively explore the environment by comparing the behavior of our curious policy to the behavior of the random policy used in prior work. To make this comparison, we execute Steps 1 and 2 visualized in Figure 6. First, our prediction model and discriminator is jointly trained on Sawyer data from the RoboNet dataset [7] by optimizing equation (7). Then, we use each policy to separately sample trajectories on a Baxter robot platform. Our curious policy was able to visit a more diverse array of states and grasp more objects than the existing random policy.

Our policy is able to significantly increase the quantity of objects that the robot grasps. Since the interaction between the robot and objects are some of the most difficult things to predict in the tabletop manipulation setting, having more data about robot-object interactions makes a collected dataset more effective in training prediction models. Figure 8 shows a histogram of when the grippers of the robot experienced non-zero forces during data collection for both the curious and the random policies. Non-zero forces indicate that an object is between the grippers, preventing them from fully closing. When following the curious policy, the robot spends a larger portion of its time grasping objects.

Our curious policy also explores a different distribution of robot states. Figure 7 shows a heatmap of the amount of time the robot's end effector spends at each location in the xy-plane. The curious policy explores more interesting regions of the state space, such as edges of the bins. The walls of the bin are interesting because they block the motion of objects, causing the objects to have more complicated dynamics than when they are in the center of the bin.

B. Prediction Performance Improvement

We also evaluated the ability of the model trained with data collected under different policies to perform prediction

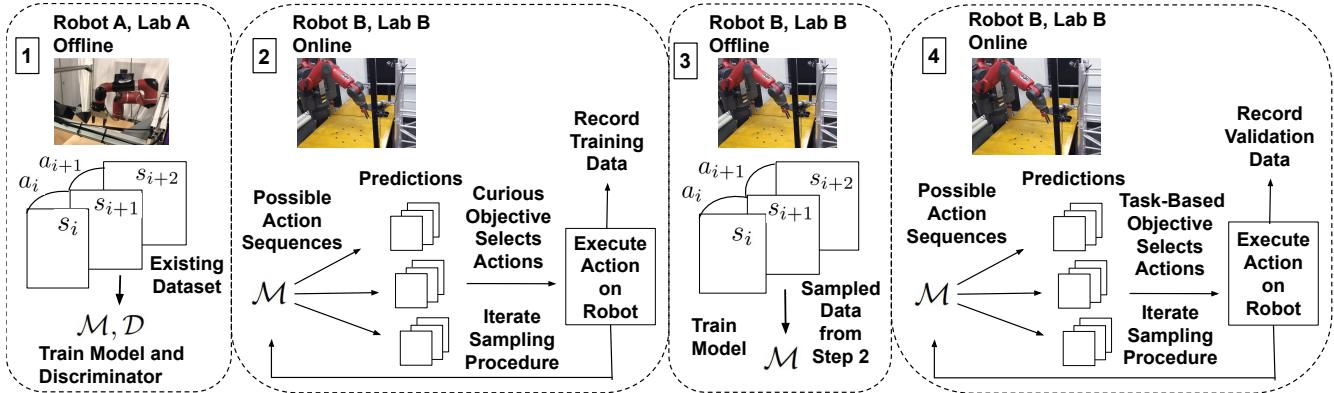


Fig. 6: The process used for online training with a curiosity objective provided by the loss from our discriminator network in a domain transfer problem. The model and the discriminator are initially trained on an existing dataset from domain A (1). The model and discriminator are used to select and execute sequences of actions that optimize the curiosity objective in domain B, generating a new dataset (2). The dataset from domain B is used to train the model (3). The model is used to select sequences of actions that optimize a task-based objective, allowing the robot to perform useful tasks in domain B (4).

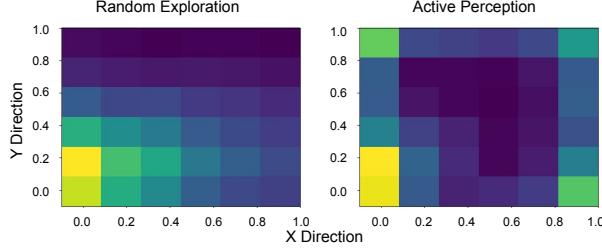


Fig. 7: Heat map of the state space regions explored by each policy over 650 trajectories. Regions which are more yellow indicate a higher count for the end effector of the Baxter arm accessing that discretized x-y region. Our curious model explores all corners of the state space, focusing on the edges where objects accumulate, while the random exploration remains near its starting location.

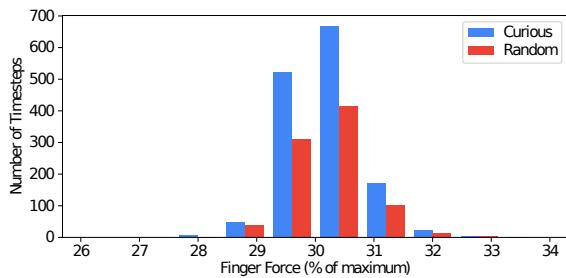


Fig. 8: Histograms of the non-zero gripper forces experienced while executing each policy. Each policy was executed for 650 trajectories of 30 timesteps. Non-zero force occurs when a large enough object is grasped by the robot's fingers. The curious policy spends significantly more time grasping objects than the random policy.

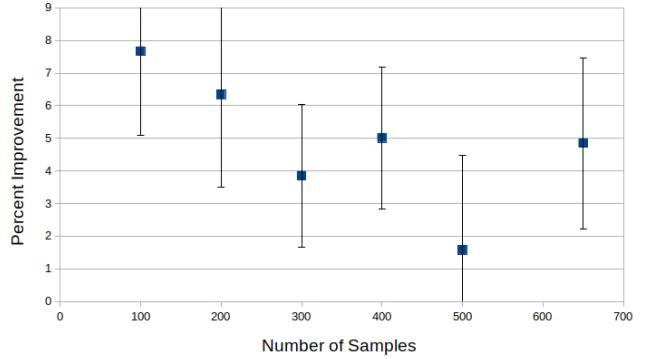


Fig. 9: The blue squares show mean percent improvement in L_2 error for the prediction model trained with curious data over the prediction model trained with the random data on the control dataset. The bars show standard error of the percent improvement. The prediction model trained with curious data performs better by more than the standard error on all but one quantity of samples. The improved performance is particularly pronounced when lower numbers of samples are used in training.

for robotic control by executing Steps 3 and 4 visualized in Figure 6.

The samples collected with our curious policy enable statistically significant prediction improvement on control tasks over samples collected with our random policy. Figure 9 shows the L_2 error improvement for the model trained with the data collected with the curious policy at different numbers of samples. The blue squares in the figure show the mean percent improvement in L_2 error, and the bars show the standard error of the percent improvement. The improvement provided by our method is statistically significant when the entire range of the error bars in Figure 9 is positive. Error improvement for the curious policy is especially pronounced at lower numbers of samples. Qualitative prediction results

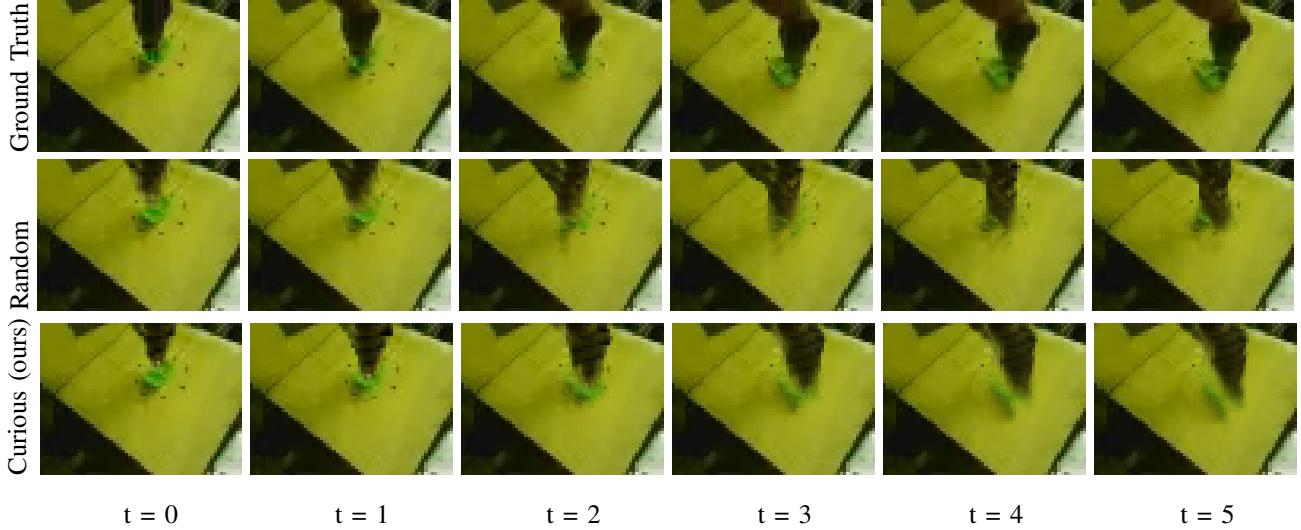


Fig. 10: Example predictions on the control dataset. All models were trained with 650 samples. In the model trained with the curious data, the object becomes more blurry, but its motion is much more accurate.

are shown in Figure 10. The model trained with the curious data more accurately tracks the position of the object.

Given an improved prediction model, we experimented with planning performance. We used CEM for planning with a task-based objective function first proposed by Finn and Levine [2], using the exact formulation and implementation from Dasari et al. [7]. We executed the same pushing tasks presented in Dasari et al. [7]. Despite the significant improvement in our prediction model, we found that the models trained with curious and random data resulted comparable control performance. Experimentally, prior work has shown an association between improved prediction performance and improved performance in control in this prediction-planning pipeline [3], [9]. However, theoretical results in model predictive control do not guarantee that improved prediction performance will result in improved control performance for pipelines using video prediction models with CEM planners. In future work, we intend to investigate the conditions in which improved prediction performance yields improved control performance for this pipeline in our experiments.

VI. CONCLUSION

We presented an adversarial curiosity approach which we use to actively sample data used to train a prediction model. Our method optimizes an objective given by the score from a discriminator network to choose the sequence of actions that corresponds to the least realistic sequence of predicted observations. We compared our approach in simulated task execution to existing model-based curiosity methods and demonstrated significant advantages of our method in the case of compute constrained environments. We then integrated our exploration strategy in a practical robotic manipulation pipeline in which existing model-based curiosity methods could not run due to computational intractability. We demonstrated increased sample efficiency and improved prediction performance in this setting for a domain

transfer problem. In future work, we plan to look into ways to guarantee improvement with active exploration methods for manipulation task performance in robotic learning pipelines.

ACKNOWLEDGEMENTS

The authors are grateful for support through the Curious Minded Machines project funded by the Honda Research Institute and the grant NSF-CPS-2038873. We would also like to thank Kenneth Chaney for compute support.

REFERENCES

[1] C. Finn, I. Goodfellow, and S. Levine, “Unsupervised learning for physical interaction through video prediction,” *NIPS*, 2016.

[2] C. Finn and S. Levine, “Deep Visual Foresight for Planning Robot Motion,” *ICRA*, 2017.

[3] F. Ebert, C. Finn, A. X. Lee, and S. Levine, “Self-supervised visual planning with temporal skip connections,” *arXiv preprint arXiv:1710.05268*, 2017.

[4] F. Ebert, S. Dasari, A. X. Lee, S. Levine, and C. Finn, “Robustness via retrying: Closed-loop robotic manipulation with self-supervised learning,” *arXiv preprint arXiv:1810.03043*, 2018.

[5] F. Ebert, C. Finn, S. Dasari, A. Xie, A. Lee, and S. Levine, “Visual Foresight: Model-Based Deep Reinforcement Learning for Vision-Based Robotic Control,” *arXiv preprint*, dec 2018.

[6] M. Zhang, S. Vikram, L. Smith, P. Abbeel, M. J. Johnson, and S. Levine, “SOLAR: Deep Structured Representations for Model-Based Reinforcement Learning,” *International Conference on Machine Learning*, aug 2018.

[7] S. Dasari, F. Ebert, S. Tian, S. Nair, B. Bucher, S. Singh, K. Schmeckpeper, S. Levine, and C. Finn, “Robonet: Large-scale multi-robot learning,” *CoRL*, 2019.

[8] A. Xie, F. Ebert, S. Levine, and C. Finn, “Improvisation through Physical Understanding: Using Novel Objects as Tools with Visual Foresight,” *RSS*, 2019.

[9] K. Schmeckpeper, A. Xie, O. Rybkin, S. Tian, K. Daniilidis, S. Levine, and C. Finn, “Learning predictive models from observation and interaction,” *arXiv preprint arXiv:1912.12773*, 2019.

[10] S. Tian, F. Ebert, D. Jayaraman, M. Mudigonda, C. Finn, R. Calandra, and S. Levine, “Manipulation by feel: Touch-based control with deep predictive models,” *ICRA*, 2019.

[11] A. Byravan, F. Leeb, F. Meier, and D. Fox, “Se3-pose-nets: Structured deep dynamics models for visuomotor planning and control,” *ICRA*, 2017.

[12] D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson, “Learning latent dynamics for planning from pixels,” *ICML*, 2019.

[13] L. Kaiser, M. Babaeizadeh, P. Milos, B. Osinski, R. H. Campbell, K. Czechowski, D. Erhan, C. Finn, P. Kozakowski, S. Levine, R. Sepassi, G. Tucker, and H. Michalewski, “Model-based reinforcement learning for atari,” 2019.

[14] P. Shyam, W. Jaśkowski, and F. Gomez, “Model-based active exploration,” *ICML*, 2019.

[15] D. Pathak, D. Gandhi, and A. Gupta, “Self-Supervised Exploration via Disagreement,” *ICML*, 2019.

[16] J. Schmidhuber, “Formal theory of creativity, fun, and intrinsic motivation (1990;2013;2010),” *IEEE Trans. on Auton. Ment. Dev.*, vol. 2, no. 3, pp. 230–247, 2010.

[17] ———, “A possibility for implementing curiosity and boredom in model-building neural controllers,” in *Proc. of the First Int. Conf. on Simul. of Adaptive Behavior*, 1990.

[18] A. Jaegle, V. Mehrpour, and N. Rust, “Visual novelty, curiosity, and intrinsic reward in machine learning and the brain,” *arXiv preprint arXiv:1901.02478*, 2019.

[19] N. Roy and A. McCallum, “Toward optimal active learning through sampling estimation of error reduction,” *ICML*, 2001.

[20] D. A. Cohn, Z. Ghahramani, and M. I. Jordan, “Active learning with statistical models,” *NIPS*, 1996.

[21] R. Bajcsy and M. Campos, “Active and exploratory perception,” *CVGIP: Image Understanding*, vol. 56, no. 1, pp. 31–40, 1992.

[22] M. Bellemare, S. Srinivasan, G. Ostrovski, T. Schaul, D. Saxton, and R. Munos, “Unifying count-based exploration and intrinsic motivation,” *NIPS*, 2016.

[23] G. Ostrovski, M. G. Bellemare, A. van den Oord, and R. Munos, “Count-based exploration with neural density models,” *ICML*, 2017.

[24] L. Lee, B. Eysenbach, E. Parisotto, E. Xing, S. Levine, and R. Salakhutdinov, “Efficient exploration via state marginal matching,” *arXiv preprint arXiv:1906.05274*, 2019.

[25] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell, “Curiosity-driven exploration by self-supervised prediction,” *CVPR*, 2017.

[26] Y. Burda, H. Edwards, A. Storkey, and O. Klimov, “Exploration by random network distillation,” *arXiv preprint arXiv:1810.12894*, 2018.

[27] B. Bucher, A. Arapin, R. Sekar, F. Duan, M. Badger, K. Daniilidis, and O. Rybkin, “Perception-driven curiosity with bayesian surprise,” *RSS Workshop on Combining Learning and Reasoning for Human-Level Robot Intelligence*, 2019.

[28] B. C. Stadie, S. Levine, and P. Abbeel, “Incentivizing exploration in reinforcement learning with deep predictive models,” *arXiv preprint arXiv:1507.00814*, 2015.

[29] I. Osband, C. Blundell, A. Pritzel, and B. Van Roy, “Deep exploration via bootstrapped dqn,” *NIPS*, 2016.

[30] R. Houthooft, X. Chen, Y. Duan, J. Schulman, F. De Turck, and P. Abbeel, “Vime: Variational information maximizing exploration,” *NIPS*, 2016.

[31] M. Kearns and S. Singh, “Near-optimal reinforcement learning in polynomial time,” *Machine Learning*, vol. 49, no. 2-3, pp. 209–232, 2002.

[32] M. Lopes, T. Lang, M. Toussaint, and P.-Y. Oudeyer, “Exploration in model-based reinforcement learning by empirically estimating learning progress,” *NIPS*, 2012.

[33] R. Sekar, O. Rybkin, K. Daniilidis, P. Abbeel, D. Hafner, and D. Pathak, “Planning to explore via self-supervised world models,” *ICML*, 2020.

[34] S. Bechtle, A. Rai, Y. Lin, L. Righetti, and F. Meier, “Curious ilqr: Resolving uncertainty in model-based rl,” *ICML Workshop on RL4RealLife*, 2019.

[35] J. Schmidhuber, “Making the world differentiable: On using self-supervised fully recurrent neural networks for dynamic reinforcement learning and planning in non-stationary environments,” Tech. Rep., 1990.

[36] ———, “Unsupervised minimax: Adversarial curiosity, generative adversarial networks, and predictability minimization,” *CoRR*, vol. abs/1906.04493, 2019.

[37] B. Settles, “Active learning literature survey,” University of Wisconsin–Madison, Computer Sciences Technical Report 1648, 2009.

[38] D. D. Lewis and W. A. Gale, “A sequential algorithm for training text classifiers,” *SIGIR*, 1994.

[39] A. J. Joshi, F. Porikli, and N. Papanikolopoulos, “Multi-class active learning for image classification,” *CVPR*, 2009.

[40] Y. Yang, Z. Ma, F. Nie, X. Chang, and A. G. Hauptmann, “Multi-class active learning by uncertainty sampling with diversity maximization,” *IJCV*, vol. 113, no. 2, pp. 113–127, 2015.

[41] H. S. Seung, M. Opper, and H. Sompolinsky, “Query by committee,” *COLT*, 1992.

[42] D. Cohn, L. Atlas, and R. Ladner, “Improving Generalization with Active Learning,” *Machine Learning*, vol. 15, no. 2, pp. 201–221, 1994.

[43] P. Melville and R. J. Mooney, “Diverse ensembles for active learning,” *ICML*, 2004.

[44] N. Houlsby, F. Huszár, Z. Ghahramani, and M. Lengyel, “Bayesian active learning for classification and preference learning,” *arXiv preprint arXiv:1112.5745*, 2011.

[45] B. Settles, M. Craven, and S. Ray, “Multiple-instance active learning,” *NIPS*, 2008.

[46] R. Sznitman and B. Jedynak, “Active testing for face detection and localization,” *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 32, no. 10, pp. 1914–1920, 2010.

[47] A. Vezhnevets, V. Ferrari, and J. M. Buhmann, “Weakly supervised structured output learning for semantic segmentation,” *CVPR*, 2012.

[48] R. Moskovitch, N. Nissim, D. Stodel, C. Feher, R. Englert, and Y. Elovici, “Improving the detection of unknown computer worms activity using active learning,” *AAAI*, 2007.

[49] K. Konyushkova, R. Sznitman, and P. Fua, “Learning active learning from data,” *NIPS*, 2017.

[50] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, “Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor,” *arXiv preprint arXiv:1801.01290*, 2018.

[51] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, “Openai gym,” *arXiv preprint arXiv:1606.01540*, 2016.

[52] A. X. Lee, R. Zhang, F. Ebert, P. Abbeel, C. Finn, and S. Levine, “Stochastic Adversarial Video Prediction,” *arXiv preprint*, apr 2018.

[53] R. Rubinstein, “The Cross-Entropy Method for Combinatorial and Continuous Optimization,” *Methodology And Computing In Applied Probability*, vol. 1, no. 2, pp. 127–190, 1999.