

Learning Accurate Long-term Dynamics for Model-based Reinforcement Learning

Nathan Lambert¹, Albert Wilcox¹, Howard Zhang¹, Kristofer S. J. Pister¹, and Roberto Calandra²

Abstract—Accurately predicting the dynamics of robotic systems is crucial for model-based control and reinforcement learning. The most common way to estimate dynamics is by fitting a one-step ahead prediction model and using it to recursively propagate the predicted state distribution over long horizons. Unfortunately, this approach is known to compound even small prediction errors, making long-term predictions inaccurate. In this paper, we propose a new parametrization to supervised learning on state-action data to stably predict at longer horizons – that we call a trajectory-based model. This trajectory-based model takes an initial state, a future time index, and control parameters as inputs, and directly predicts the state at the future time index. Experimental results in simulated and real-world robotic tasks show that trajectory-based models yield significantly more accurate long term predictions, improved sample efficiency, and the ability to predict task reward. With these improved prediction properties, we conclude with a demonstration of methods for using the trajectory-based model for control.

I. INTRODUCTION

Model-based reinforcement learning (MBRL) has emerged as a compelling approach for control across multiple domains of robotics by using a learned forward dynamics model for planning [1], [2], [3]. At the base of most current MBRL algorithms is the idea of learning *one-step ahead forward dynamics models* – that is, given a state and an applied action the model will predict the future state. Once a one-step ahead model is learned, it can subsequently be used for planning or control by recursively applying the model to the state to unroll long-horizon trajectories. While this approach has proved successful in many domains and is well-motivated from the related field of optimal control [4], [5], it is well understood that a major limit of current MBRL algorithms is the insufficient accuracy of the predictions over long-horizons. In fact, one-step models critically suffer from compounding errors of the dynamics [2], [6], [7], and recent work has shown that the performance of one-step ahead prediction models is not necessarily correlated with downstream control performance [8].

In this work, we propose and study a new class of feed-forward dynamics models focused on capturing not the behavior of single steps, but the long-term time dependant evolution of a trajectory as a whole. The main intuitions behind this model are that: 1) sequences of actions and states are usually strongly correlated with their neighbors across time (i.e., across a trajectory) and 2) in quantifying

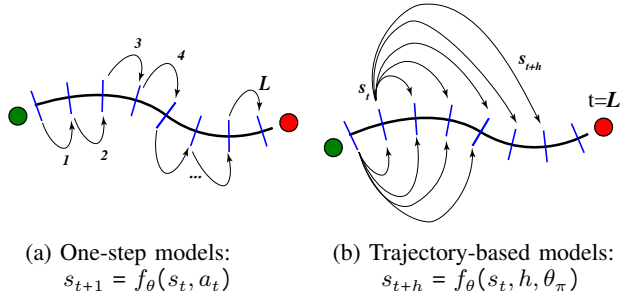


Fig. 1: The model formulation for one-step models (a) and our new **trajectory-based models** (b).

the quality of long-horizon predictions, it might be preferable to have higher uncertainty over the single steps, if the prediction of the trajectory as a whole is more accurate. This is particularly crucial when the dynamics models are used for planning, where the relative ranking of the trajectories will directly impact the decision making (e.g., it might be relatively unimportant to know how we reach a location as long as we know that we can reach it accurately).

Our contribution is to propose a new class of deep predictive models, **trajectory-based models**, that are focused on capturing the long-term, time-dependant nature of dynamics. The new models re-frame the supervised learning problem to predict future states from the parameters of a controller parameters and a time index, rather than actions. The trajectory-based prediction steps is compared visually to individual steps in Fig. 1.

This provides several benefits compared to traditional one-steps ahead models: improved data-efficiency, better computational complexity which can be further parallelized, a more principled treatment of uncertainty over trajectories, the capability of modeling continuous time steps, and finally more accurate long-term predictions. We evaluate our trajectory-based approach on simulated and real-world data, and demonstrate the advantages of this approach for tasks with a long-horizon to: a) accurately predict long-term behaviors of stable robot dynamics, and clearly outperforms the one-step models after about 50 time-steps; b) improve the sample efficiency; c) accurately predict downstream task reward. With these benefits, we show how the model can be used in an iterative learning task and planning via model predictive control. We believe that this approach is a small but significant step towards understanding and overcoming the limitations of the dynamics models in current model-based reinforcement learning algorithms.

Corresponding author: Nathan O. Lambert, nol@berkeley.edu

¹Department of Electrical Engineering and Computer Sciences, University of California, Berkeley, USA.

²Facebook AI Research, Menlo Park, CA, USA.

II. RELATED WORKS

A. Model Learning for Control

The foundation of trajectory-based models is the context of using predictions for control [5] and learning dynamics for unknown systems [9]. In robotics, combining these two directions is the method of data-driven model predictive control (MPC), which directly leverages a dynamics model to choose an action [10], [11], [12]. In recent years, data-driven approaches for model learning and control have become more prevalent. Model-based reinforcement learning (MBRL) is the iterative framework of an agent acting to collect data, modeling the transition function with said data, and leveraging the model for control. MBRL has shown impressive performance and sample efficiency on many robotics tasks [13], [1], [2], [14]. Model-based planning approaches have been successfully using one-step predictions and model predictive control (MPC) to select best actions [1], [15], [16]. This paper proposes a new technique of predicting trajectories for data-driven modeling of unknown dynamics in hopes to improve robotic learning tasks in MBRL, especially with MPC. System identification (SI) has a large history for learning models for control, but generally is not formulated as an iterative learning problem [17].

B. Predicting with Single Steps Ahead

Single-step dynamics models are effective across many domains, but they suffer from compounding errors. [2] showed excellent downstream control performance with probabilistic neural network (NN) ensembles, and similar approaches have been applied to many robotic tasks [1], [16]. Methods mitigating the compounding error of multi-step predictions include re-framing the prediction in the lens of imitation learning [18], using a multi-step value prediction model [6], and leveraging a model with flexible prediction horizon [7]. Additionally, single-step Gaussian Processes (GPs) have been applied to multiple lower-dimensional robotic learning tasks [13]. GP models have been studied in many contexts, such as the long term predictions of pedestrians [19].

C. Predicting Trajectories

Predicting long-term trajectories has a long-history, dating back to autoregressive-moving-average models (AR-MAX) [9], and has been a growing area of research as robotic complexity outgrew the performance of historical approaches. A common approach is to modify one-step prediction training to account for correlation across a trajectory with NNs or GPs. Specifically, [20] proposes a time-weighted loss function to learn unstable state-space systems, but does not apply it to high-dimensional, nonlinear systems. [8] showed increased correlation between model-accuracy and downstream reward when training with batches sampled from a single-trajectory instead of random transitions, but predictions still suffer from compounding errors.

Other related methods have attempted to create direct, structural links between models and trajectory training data. With high dimensional images (rather than states) [21] uses an auto-regressive, recurrent network to predict observations

in a latent state-space. [22] proposes a multi-step Gaussian Process for learning robotic control and the approach is studied further using the correlation between prediction steps in [23]. [24] suggests using a new kernel in GPs to correlate across trajectory data and using simulated model rollouts to predict reward as a prior GP mean function. These methods all leverage one-step models to create long-term dynamics, which we differentiate from by including time-dependence. Other exciting avenues of long-term prediction of dynamical systems are with neural ordinary differential equations [25] and long short-term memory blocks (LSTMs) [26], but both are yet to be successfully applied to proprioceptive (*i.e.*, not vision based) robotics control tasks. We baseline our method against LSTMs which have implicit time dependence, where our method includes it explicitly in the input.

III. BACKGROUND

Markov Decision Process In this work, we formulate the prediction problem in the context of a Markov Decision Process (MDP) [27]. An MDP is defined by the state $s_t \in \mathbb{R}^{d_s}$, the action $a_t \in \mathbb{R}^{d_a}$, the reward function $r(s_t, a_t) \in \mathbb{R}$, and the transition function $f : \mathbb{R}^{d_s \times d_a} \mapsto \mathbb{R}^{d_s}$. In this work, we look to learn a parametrized dynamics model $f_\theta(s_t, a_t)$ to approximate the true transition function. We consider agents utilizing a control law $\pi(\cdot)$ to act in the environment for repeated trajectories τ_k . Aggregating the tuples of experience, the agent receives a dataset $\mathcal{D} = \{(s_i, a_i, s_{i+1})\}_{i=1}^N$.

One-step Dynamics Models Given the state s_t , and the action a_t , one-step dynamics models predicts the next state s_{t+1} (or often a delta compared to the current state) as

$$s_{t+1} = f_\theta(s_t, a_t). \quad (1)$$

For the prediction of trajectories the dynamic model is then recursively applied such that

$$s_{t+h} = f_\theta(\dots f_\theta(f_\theta(s_t, a_t), a_{t+1}) \dots, a_{t+h}). \quad (2)$$

For any model used to represent f (including deep neural networks), there will inevitably be a prediction error $\epsilon_t = \|\hat{s}_t - s_t\|$. Unfortunately, this error is not well studied can grow in a multiplicative manner with each recursion of the function evaluation as

$$s_{t+h} = f_\theta(\dots f_\theta(f_\theta(s_t, a_t) + \epsilon_t, a_{t+1}) + \epsilon_{t+1} \dots, a_{t+h}) + \epsilon_{t+h}. \quad (3)$$

Hence, even small numerical errors will ultimately make predictions over long trajectories inaccurate. This error growth with one-step models is often referred to as the compounding error problem in MBRL.

Recurrent Dynamics Models (RNN) RNN models predict the next state, s_{t+1} , as a function of the current hidden state of the network, $s_{t+1} = f_\theta(\zeta_t)$. The hidden state is outputted from the recurrent nodes of the network: given the current hidden state, ζ_t , state, s_t , and action, a_t , the network passes information, as $\zeta_{t+1} = g_\theta(s_t, \zeta_t, a_t)$. In this paper, we study the ability of LSTMs, which are often used to mitigate the vanishing gradient problem [26], but they are also difficult to train in practice.

IV. TRAJECTORY-BASED DYNAMICS MODELS

A. Prediction Formulation

We now describe our new trajectory-based dynamics models (which we refer in the rest of the manuscript as T) which focus on modeling trajectories over time rather than individual steps. There are two main intuitions behind the adoption of this type of model: 1) for control purposes it is often more valuable to have an accurate overall trajectory prediction compared to accurately predict single steps (which might compound error over long-term). This is even more important when planning, since for planning the relative ranking of the trajectories is what determines the eventual actions applied by control scheme such as MPC. 2) trajectories are often strongly correlated in space and time; however, single-step models do not have efficient mechanism to enforce that.

To address the error compounding, an idea would be to replace the recursive call of Eq. (2) with a n^{th} step prediction

$$s_{t+h} = f_{\theta}(s_t, a_t, a_{t+1}, \dots, a_{t+h}), \quad (4)$$

which does not require recursion, and is thus more likely to produce stable long-term predictions. However, here we can observe how the dimensionality of the model to be learned depends on the length of the prediction into the future h , and that in addition, this model is generally only capable of predicting the resulting n^{th} step ahead prediction, but not its intermediate steps (this is not true for RNNs, but their formulation is more similar to Eq. (2)). A first variant of this n^{th} step ahead formulation would be to observe that the sequence of action $a_t, a_{t+1}, \dots, a_{t+h}$ is typically generated by a generic, but single controller $\pi(\cdot)$ determined by parameters θ_{π} , and thus we can rewrite as

$$s_{t+h} = f_{\theta}(s_t, \theta_{\pi}). \quad (5)$$

As long as the dimensionality of θ_{π} is smaller than the dimensionality of $a_t, a_{t+1}, \dots, a_{t+h}$, this would result in an effective reduction of the dimensionality of our dynamics model and thus improved data-efficiency. However, once again this model only allows us to predict the final state but not the trajectory that led us there. Adding the notion of dynamic time-prediction is the final conceptual change to attempt to accurately predict long-term system dynamics in a data-efficient manner, by which we index the starting state at time t and directly predict to the future, variable horizon of h steps with one forward pass. The trajectory-based models predict the evolution from a starting state s_t , subject to the control parameters θ_{π} , to a future-time index $t+h$, as

$$s_{t+h} = f_{\theta}(s_t, h, \theta_{\pi}). \quad (6)$$

Compared to the traditional recursive one-step ahead formulation of Eq. (1), this formulation provides several benefits that we now detail.

B. Benefits of Trajectory-based Models

Data-efficiency One advantage of this formulation is that we can perform a re-labeling trick over the dataset of collected trajectories, to significantly augment the dataset used

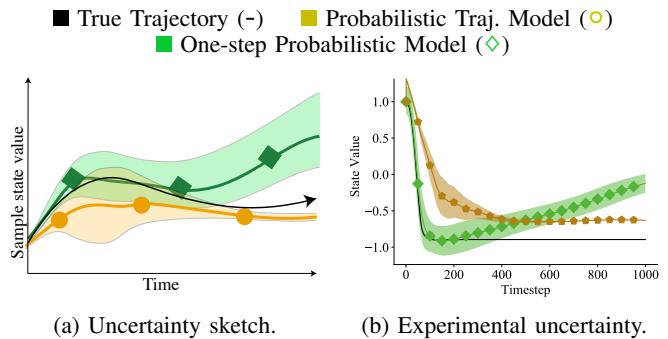


Fig. 2: The trajectory-based models have prediction uncertainty proportional to the epistemic uncertainty in the training dataset. (Left) a sketch of the uncertainty mechanism. Trajectory-based models have uncertainty that can shrink when more confident in the dynamics, which is opposed to one-step models that have predicted uncertainties that diverge at long horizons. (Right) an example trial of a robotic prediction (from the reacher task in Sec. V-A) highlights this uncertainty propagation with a probabilistic trajectory-based model and an one-step probabilistic model (P).

to train the trajectory-based model. We assume a dataset \mathcal{D} of n collected trajectories $\mathcal{D} = \{\tau^n\}$, each of fixed length L . For each collected trajectory $\tau^j = [s_0, \dots, s_L]$ we can now extract $L-1$ subtrajectories $\tau_i^j = [s_i, \dots, s_L]$ for $i = 0 \dots L-1$, and use them all for training the trajectory-based model. By training on all sub-trajectories, the model gains two strengths: 1) it can predict into the future from any state, not just those given as initial states from an environment, and 2) the number of training points grows proportional to the square of trajectory length, as

$$N_{\text{train}} = n \sum_{t=1}^L t = n \frac{(L)(L-1)}{2} \approx nL^2. \quad (7)$$

This results in models that better exploit the temporal structure of the systems, while using less data.

Computationally Efficient Planning The trajectory-based models have a useful property of directly predicting entire trajectories instead of imagined roll-outs composed of repeated model evaluations. For prediction propagation, by only passing in a vector of time horizons h from a current state, a planner can evaluate the future with one forward pass, alleviating the computational burden (as well as the compounding, multiplicative error) associated with evaluating sequentially many steps of one-step models. In our model the predictions in a trajectory do not depend on the prediction at the previous step, which can dramatically increase the control frequency when planning online.

Capturing Empirical Distribution over Trajectories One-step models commonly suffer from the issue of uncertainty explosion, where the predicted uncertainty over a trajectory typically keep increasing, and does not match the empirical uncertainty from the data. By propagating time directly, our probabilistic trajectory-based model can

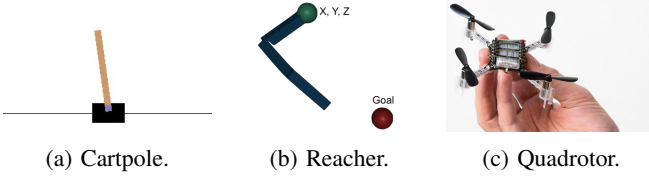


Fig. 3: Experimental platforms.

instead capture the uncertainty of variation in dynamics in the training set (*i.e.*, the model is more uncertain in areas of rapid movement and can become confident when motion converges), and the empirical uncertainty over trajectories. The uncertainty propagation is drawn in Fig. 2a and an example experiment is shown in Fig. 2b; both are compared to one-step models that have diverging uncertainty as the predicted states leave the training distribution. Stable uncertainty estimates convey promise when planning on robotic hardware, where action choices are balanced against model uncertainty due to high cost-per-test.

Continuous Time Traditional one-step ahead models assume a discrete quantization of time such that the sampling frequency is constant. Instead, our model is agnostic to the use of discrete or continuous time, since the model can make use of data collected at arbitrary h and explicitly interpolate between them. While this property is not employed in the following experiments, this is a very desirable property that we aim to explore in future work.

V. EXPERIMENTAL RESULTS

We now evaluate the proposed trajectory-based models. In particular, we investigate the long term prediction accuracy, the ability for the trajectory-based model to predict unstable or periodic data, the sample efficiency benefit of the new parameterization, and using the new model for predicting experimental reward. Code and an expanded manuscript are available on the website: <https://sites.google.com/view/trajectory-prediction/>.

A. Experimental Setting

Model Training Using the same notation and model training formulations in [2], we use four model types: D, P, DE, PE . The deterministic model, D , and deterministic ensemble, DE , minimize the mean squared error (MSE) of predictions. The probabilistic model, P , and probabilistic ensemble, PE , minimize the negative log likelihood (NLL) of a Gaussian distribution of state transitions.

All models normalize the input states and actions to a standard normal distribution $\mathcal{N}(0, 1)$ in each dimension, and bounded control parameters are mapped to $[-1, 1]$. The feedforward models have two hidden layers of width 250, are optimized with Adam [28], with batch sizes of 32 for D, P and 64 for T , and learning rates of 2.5×10^{-5} for P models, 5×10^{-5} for D and 8×10^{-4} for T . Due to the rapid accruing of labeled data for the trajectory-based models, we cap the training set size at 1×10^5 by random downsampling. The LSTMs are trained with Adam with a learning rate of 0.1,

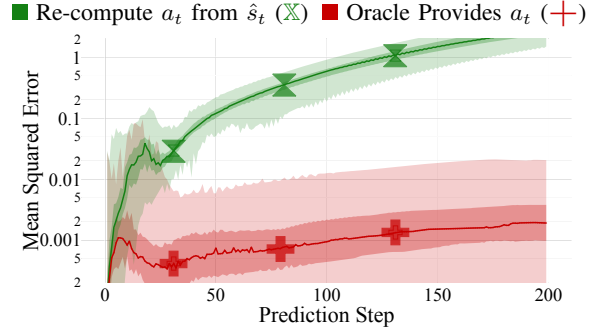


Fig. 4: The prediction accuracy of a PE on Cartpole with and without re-computing the action at each step ($N_{\text{traj}} = 100$) shows the increased error from re-computing actions.

with batches of sequences matching the trajectory length, L , and with normalization following [29].

Cartpole (Simulated) We evaluate our models thoroughly on the cartpole task where the goal is to balance a mass over a sliding cart ($d_s = 4, d_a = 1$), shown in Fig. 3a. To have a continuous reward for prediction, we introduce a new reward function $r(s_t, a_t) = -(x^2 + \theta^2)$ and the remaining details are outlined in [2]. We evaluate predictions of state and reward of cartpole agents conditioned on a Linear Quadratic Regulator (LQR) control policy. LQR solves the optimization $\min_u J(u) = \int_0^\infty s^\top Q s + a^\top R a \, dt$ for the dynamics $\dot{s} = \tilde{A}s + \tilde{B}u$. LQR control minimizes the expected cost based on a linearized dynamical system, \tilde{A}, \tilde{B} . The control policy, $\pi(\cdot)$, of LQR takes the form of state feedback, $a = -Ks$, where $K \in \mathbb{R}^{d_a \times d_s}$. For all experiments we use the following cost matrices to generate an optimal controller, K^* : $Q = \text{diag}(.5, .05, 1, .05)$, $R = [1]$, and then sample a random vector $m \in \mathbb{R}^{4 \times 1}$ uniformly from the interval $[0.5, 1.5]$ to create a variety of controllers $K_i = m_i \cdot K^* \forall i$.

Reacher (Simulated) For a higher dimensional task, we examine the 5 joint, three-dimensional, reacher manipulation task ($d_s = 15, d_a = 5$) in the Mujoco, OpenAI Gym environment [30], [31], shown in Fig. 3b. The task associated with the environment is to maneuver the end-effector of the arm from an initial position state to an end position state. To create a diverse set of data for prediction, our experiments control the agent using a Proportional-Integral-Derivative (PID) controller with randomly generated parameter vectors $K \in \mathbb{R}^{15}$. The parameters of a PID control are defined by a vector of joint angle targets, $z_d \in \mathbb{R}^5$, proportional constants, $K_p \in \mathbb{R}^5$, integrative constants, $K_I \in \mathbb{R}^5$, and derivative constants, $K_D \in \mathbb{R}^5$, for each rotatory joint. We set $K_I = 0$ for all experiments. Given the joint angle z_i and the current error $e_i = z_i - z_d$, the control command at the i^{th} joint is $u_i = K_p \cdot e_i + K_D \cdot \dot{e}_i$.

Quadrotor (Simulated & Real Hardware) We validated the prediction accuracy of trajectory-based models on a simulated and experimental low-level attitude control of a quadrotor ($d_s = 5, d_a = 4$). The quadrotor model is based off the Crazyflie [32], shown in Fig. 3c, an 27g,

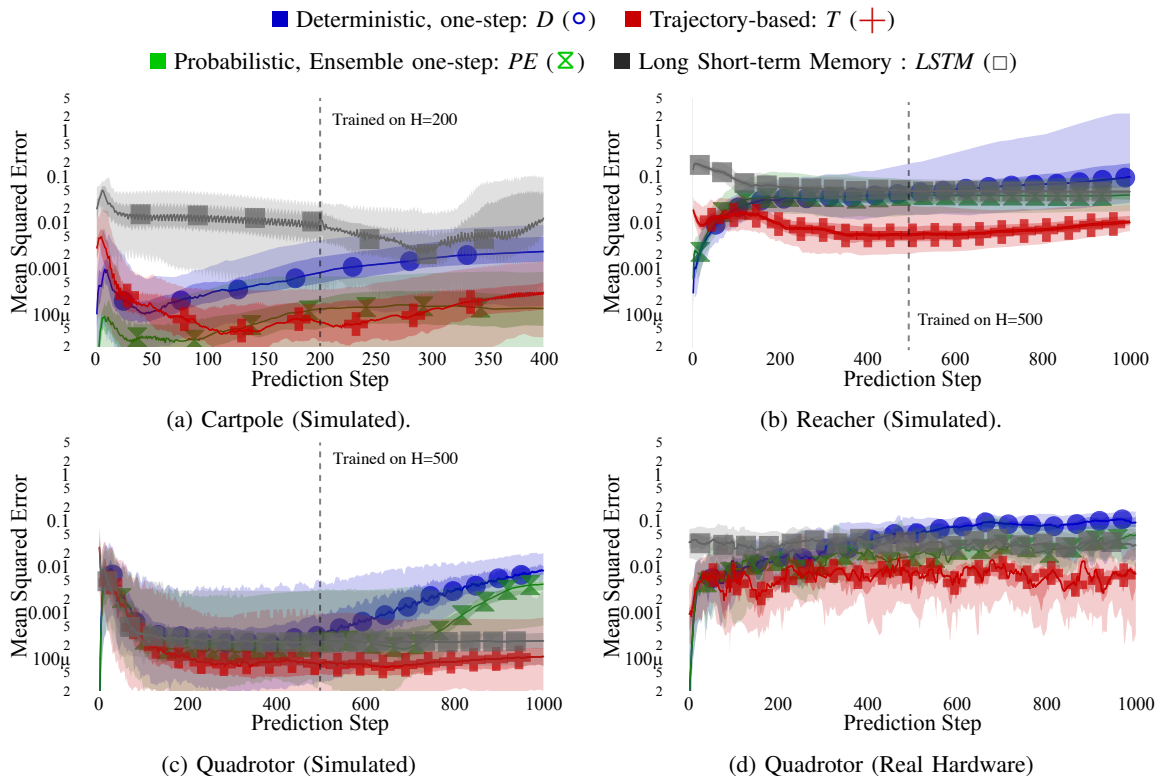


Fig. 5: The log-scale, mean-squared prediction error for the tested environments. The simulated environments are tested on 250 validation trajectories and the experimental data is tested on 16 training trajectories due to experimental restrictions. Highlighted is the median error with the 65th and 95th percentile of errors. These figures highlight two properties: 1) 5× to 10× gain in long term prediction accuracy via trajectory-based models for $h > 50$ and 2) the uncertainty in one-step prediction continues to grow with the prediction step while trajectory-based error remains stable. The vertical line indicates the length of trajectories in the training distribution.

open-source micro-aerial vehicle. The 12 state Euler-step simulation follows [33] and has uniform Gaussian noise on all state variables sampled from $\sigma \sim \mathcal{N}(0, 0.01)$. The simulated controller is a linear, pitch and roll PD controller with randomly sampled parameters. For experimental data, we collected 180s of aggressive flight data with default PID rate-controllers. This data was broken down into trajectories of length 1000 randomly, which we used to validate our prediction mechanism.

B. Long-term Prediction Accuracy

We now demonstrate the ability of the trajectory-based models (T) to more accurately predict long-horizon robotic dynamics by measuring the mean-squared error of the predicted trajectory versus the measured state, $\sum_{t=1}^H \|\hat{s}_t - s_t\|^2$. We evaluate the ability to predict horizons of over 100 steps and trajectories longer than the original training distributions. An advantage of trajectory-based models over one-step models is that T models lack a need to be given a time-series of actions from an oracle or to compute a new action from the current state. When predicting to long horizons with one-step models, the error can compound and diverge rapidly if the predicted state is used to re-compute the action, shown in Fig. 4. For a more competitive baseline in the remainder

of our experiments, the one-step models predict the next state given the original action sequences, $\{a_t\}_{t=0}^L$, and the trajectory-based models given only θ_π .

The prediction accuracy for D , PE , $LSTM$, and T models with error 65th, 95th percentiles (tested on 100 trajectories) is shown for the cartpole trained on 100 trials of 200 time-steps, Fig. 5a, 100 reacher trials of 500 time-steps, Fig. 5b, 100 simulated quadrotor trials of 500 time-steps, Fig. 5c, and 16 experimental quadrotor trajectories of 1000 time-steps, Fig. 5d. The experimental quadrotor trajectories all have the same control parameters, which the model could use to better generalize across trajectories, showing that adding time-dependence alone can improve long-term prediction accuracy. All states are normalized to a range of $[0, 1]$ before computing the MSE to match error across different states (*i.e.*, the scale of a velocity is matched to the scale of an angle). The trajectory-based models are less accurate for short horizons ($h < 25$), but converge to an improvement of up to 10× reduction in MSE for long horizons both in simulation and experiment. In practice, it is expected that some testing trajectories will extend beyond the expected length, which we evaluate by testing on trajectories of greater length than the training set. Even with out of distribution time indices, the trajectory-based models maintain their improvement in

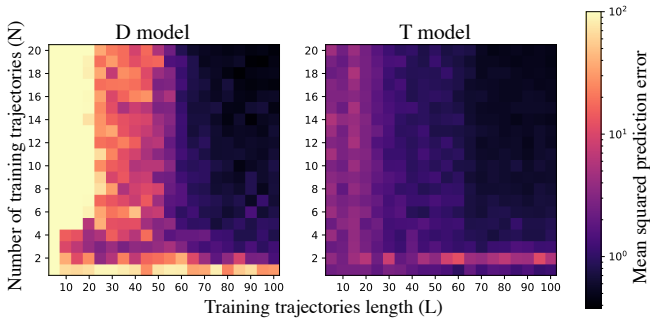


Fig. 6: The median, cumulative prediction error of 5 models sweeping over the number N and length L of reacher training trajectories given on a constant validation set of trajectories of length 100. The trajectory-based model achieves substantially better prediction accuracy with both shorter and fewer trajectories in the training set. In the context of MBRL, the trajectory based model has better *sample efficiency* by having lower cumulative prediction error when trained on fewer trajectories N (a slice at a specific length L).

accuracy over one-step models, removing any need for T models to be trained on equal-length trajectories.

C. Accelerated Data Efficiency

Having sufficient labeled data is frequently a limiting factor in deep-learning tasks. Recall from Sec. IV that the labelled data for trajectory-based models grows at a rate of the trajectory length squared, L^2 . Leveraging this accelerated accruing of data, we evaluate the ability of trajectory-based models to predict accurately in the low-data regime. In the reacher environment, we trained models on a grid of trajectory lengths, $L \in [5, 100]$, and number of trajectories, $N \in [1, 20]$, to predict a validation set of 10 trajectories of length 100. Given an initial state, s_t , we can predict a set horizon, h , into the future to obtain a simulated trajectory of states, $\hat{\tau} = \{\hat{s}_i\}_{i=t}^{t+h}$, and measure the mean-squared error (MSE) across all normalized state dimensions. In Fig. 6, the trajectory-based models show improved performance for both a) shorter trajectories (L) (another link to predicting beyond the initial training length in Sec. V-B) and b) fewer samples (N). The regime of low number of training trajectories (N) represents an area of high value to robotic tasks via its potential for reduced evaluations on a real robot.

D. Predictive Episode Reward

Predicting reward is tied to planning actions for robotic systems because if one can accurately predict rewards for an action set, then one can correctly rank actions. The trajectory-based models predict rewards in a simulated task by coupling reward prediction to stable long-term predictions. In this section we compare the predicted reward, \hat{r} , of an initial state, s_0 , and control parametrization \mathbf{K} to the true cumulative reward, $\mathbf{r} = \sum_{t=0}^L r(s_t)$.

We consider five methods for predicting the reward of an episode from an initial state and control parameters: 1) the trajectory-based models, 2) the one-step models given

the action sequence from true-states (oracle), 3) the one-step models with actions computed from predicted states (predicted-actions), 4) Gaussian Processes (GPs), and 5) neural networks (NN) predicting directly from the initial state and control parameters to sum of reward. Each candidate method is a different mapping in the following function-space: $h : (s_0, \mathbf{K}, \theta) \mapsto \mathbf{r}$, where θ carries different model formulations. The dynamics models are used to predict future states via $\hat{s}_{t+1} = f_\theta(\cdot)$, with the one-step models taking in the previous predicted state and the trajectory-based models updating the time index. The predicted trajectory reward uses the environment-defined reward function $r(s_t, a_t)$ summed over time (the reward functions are action independent). A Gaussian process (GP), defined by a mean vector, $\mu_\theta(\mathbf{x})$, and a covariance matrix, $k(\mathbf{x}_1, \mathbf{x}_2)$, or a neural network (NN) can predict the reward with no structured dynamics model. For the GP and the NN, the target rewards are normalized uniformly to $[-1, 1]$ before prediction to aid in model training.

For both the cartpole and the reacher tasks, the trajectory-based model outperforms other methods in predicted reward accuracy. The mean-squared predicted reward error across 100 trajectories is shown in Tab. I. We hypothesize that the *structured* learning of a dynamics model accurate across a trajectory improves reward prediction with knowledge of each step over predicting directly to the cumulative reward.

E. Iterative Learning of Control Parameters

We now compare the trajectory-based model against black-box optimization algorithms for the iterative learning of control parameters. At each iteration, we retrain the trajectory-based model and then use it to simulate rollouts of different control parameters. Finally, we execute the best control parameters found on the real system, and a new iteration starts. Ideally, a more accurate dynamics model over the trajectory will result in faster convergence and better performance. In our experiments, we indicate this approach as *Trajectory Optimization* and use covariance matrix adaptation evolution strategy (CMA-ES) to optimize the simulated rollouts. However, CMA-ES can be replaced with any optimizer that does not require gradients on the reward function. Trajectory Optimization generates new control parameters within the population of CMA-ES and simulates a trajectory of the trial length 200 for cartpole. The simulated reward \hat{r} is the sum over the predicted states, as in Sec. V-D. We compare this approach to the data-efficient black-box optimization algorithm Bayesian Optimization (BO) which iteratively optimize the control parameters without knowledge of the dynamics [34], [35]. The results of the experiments in Fig. 7 show that on the toy cartpole benchmark task Trajectory Optimization converge to excellent performance faster than Bayesian Optimization. This demonstrates that we can learn in an iterative manner trajectory dynamics model, and that they can successfully be applied for control.

F. Model Predictive Control with Trajectory-based Models

We now demonstrate how to use the trajectory-based model in a common control architecture – model predictive

Prediction Mechanism	Prediction Mapping	Cartpole MSE ($r - \hat{r}$) $\pm\sigma$	Reacher MSE ($r - \hat{r}$) $\pm\sigma$
Direct mapping (GP)	$(\mathbf{K}, s_0) \mapsto \hat{r}$	0.884 ± 1.873	0.127 ± 0.192
Direct mapping (NN)	$(\mathbf{K}, s_0) \mapsto \hat{r}$	0.446 ± 1.189	0.074 ± 0.112
One-step; oracle (D)	$(f_\theta(\cdot), \mathbf{a}_{t=0:L}, s_0) \mapsto \hat{r}$	0.057 ± 0.100	0.485 ± 2.048
One-step; pred. actions (D)	$(f_\theta(\cdot), \mathbf{K}, s_0) \mapsto \hat{r}$	0.979 ± 1.725	0.072 ± 0.0932
Trajectory-based (T)	$(f_\theta(\cdot), \mathbf{K}, s_0) \mapsto \hat{r}$	0.010 ± 0.033	0.005 ± 0.006

TABLE I: The mean-squared predicted reward error across 100 simulated trajectories show the strength of learning long term dynamics for predicting episode reward. The 100 trajectories have different initial states s_0 , and control parameters \mathbf{K} . The rewards are normalized per the number of trial step – 200 steps for the cartpole task and 500 for the reacher task.

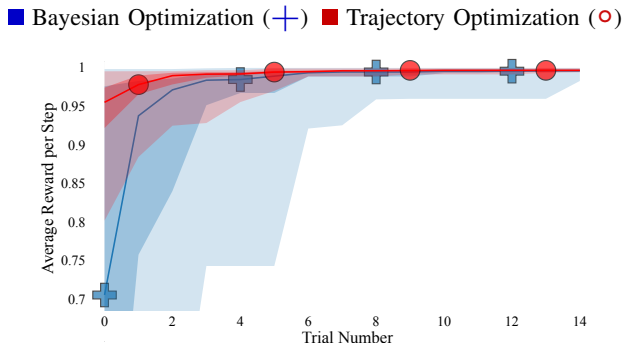


Fig. 7: The cumulative maximum reward for iterative learning on cartpole. The Trajectory Optimization with CMA-ES consistently solving the task in 2-4 trials (1 is the maximum normalized reward per step) – faster than the baseline of Bayesian Optimization. The 66th and 95th percentiles over 25 trials are shaded.

control (MPC) [10], [11]. MPC is a common tool for model-based reinforcement learning [2], [1], [15], and originated in the study of optimal control leveraging predictions to make decisions [4], [5]. MPC with learned one-step dynamics models is formulated as

$$a_t^* = \arg \max_{u_{t:t+\tau}} \sum_{i=0}^{\tau} r(\hat{s}_{t+i}, a_{t+i}), \quad \hat{s}_{t+1} = f_\theta(\hat{s}_t, a_t). \quad (8)$$

With the trajectory-based model, the control formulation needs a modification in how the candidate solutions are selected. Sampling over control policies and computing the action from the current state, the new MPC formulation is

$$\theta_{\pi,t}^* = \arg \max_{\theta_{\pi,t:t+\tau}} \sum_{i=0}^{\tau} r(\hat{s}_{t+i}, a_{t+i}), \quad (9)$$

$$\hat{s}_{t+\tau} = f_\theta(s_t, \theta_{\pi,t}, t + \tau), \quad a_t^* = \theta_{\pi}^*(t). \quad (10)$$

MPC is known to be computationally intensive – where some robots lack the computing infrastructure to run recent MBRL methods online – so we compare leveraging the trajectory-based optimization only from the first state, and running that policy through the remaining of the trial. This is a starting point, and re-planning online with varying update frequencies (holding the chosen policy for T steps) would allow flexibility in control.

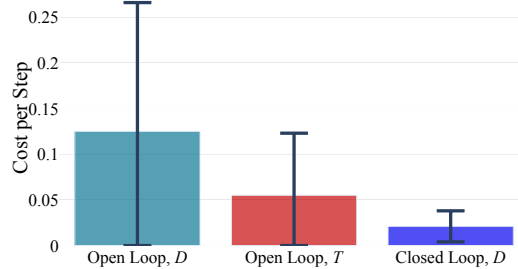


Fig. 8: Comparison between trajectory-based and one-step ahead models planning from the initial state (i.e., open-loop) and the one-step model re-planning after each step (i.e., closed loop). The trajectory-based planning improves performance compared to the one-step ahead in open-loop. However, the one-step model used in closed-loop performs the best, at the expense of increased computational complexity for re-planning. The mean and standard deviation of the reward per step in the cartpole task are reported for 25 trials.

As a comparison to iterative learning of one set of control parameters, as in Sec. V-E, we compare the performance of MPC to that if the optimization is only run on the first time-step. In this case we maintain random sampling to mirror common applications of MPC in MBRL, while the Trajectory Optimization in Sec. V-E leveraged the CMA-ES optimizer. The preliminary results for the planning methods on the cartpole task are shown in Fig. 8. The trajectory-based model is limited by aggregating labelled data with MPC because it requires the sub-trajectories that it is labelled on to have constant control parameters. This can be partially overcome by re-planning at a lower frequency, but is an important direction for future work to better integrate the new models into existing MBRL literature.

VI. DISCUSSION AND FUTURE WORK

There is *no free lunch* with respect to numerical dynamics modelling. The trajectory-based model’s accurate long-term prediction and rapid data accruing is valuable to data-driven methods with parametrized controllers. One-step models will likely remain useful for other algorithms less focused on the long-term future, such as MBPO [14], or in situations where data is non-episodic, so applying the time-dependant structure could be forced. Another limitation of the trajectory-based model relative to the entire space of

MRBL is that this model is designed for scenarios where the generalization across similar trajectories is useful, likely limiting the potential generalization of the model to one type of task instead of an entire state space. By setting training structures and input-output pairs, any dynamics modeling algorithm is prioritizing its model capacity on a certain task. Finally, our approach was demonstrated mainly on modeling low-dimensional parametrized controllers. Future work will focus on evaluating and extending this approach with high-dimensional neural network policies, as well as applying it in online model-based reinforcement learning scenarios, including real-world hardware.

VII. CONCLUSION

While one-step models have been successful across numerous robotic applications, this paper re-evaluates the paradigm of one-step models to better match the prediction mechanism with that of predicting and controlling trajectories. In this paper, we introduce a new framework for the learning of trajectory-based dynamics for iterative tasks. Trajectory-based models have some limitations, such as the need to access to policy parameters instead of just actions that can be gathered with multiple controllers. However, our trajectory-based models also present several theoretical advantages, such as data-efficiency, computational cost, parallelization, continuous time predictions, and better handling of uncertainty. The experimental results demonstrate strong performance in long-horizon prediction for simulated and real-world systems. We believe that better understanding of trajectory-based models is an important step to overcome the current limitations of MBRL.

REFERENCES

- [1] G. Williams, N. Wagener, B. Goldfain, P. Drews, J. M. Rehg, B. Boots, and E. A. Theodorou, "Information theoretic mpc for model-based reinforcement learning," in *International Conference on Robotics and Automation*, 2017, pp. 1714–1721.
- [2] K. Chua, R. Calandra, R. McAllister, and S. Levine, "Deep reinforcement learning in a handful of trials using probabilistic dynamics models," in *Neural Information Processing Systems*, 2018.
- [3] J. Schrittwieser, I. Antonoglou, T. Hubert, K. Simonyan, L. Sifre, S. Schmitt, A. Guez, E. Lockhart, D. Hassabis, T. Graepel *et al.*, "Mastering atari, go, chess and shogi by planning with a learned model," *Nature*, vol. 588, no. 7839, pp. 604–609, 2020.
- [4] C. E. Garcia, D. M. Prett, and M. Morari, "Model predictive control: theory and practice—a survey," *Automatica*, vol. 25, no. 3, 1989.
- [5] D. E. Kirk, *Optimal control theory: an introduction*. Courier Corporation, 2004.
- [6] K. Asadi, D. Misra, S. Kim, and M. L. Littman, "Combating the compounding-error problem with a multi-step model," *arXiv preprint arXiv:1905.13320*, 2019.
- [7] C. Xiao, Y. Wu, C. Ma, D. Schuurmans, and M. Müller, "Learning to combat compounding-error in model-based reinforcement learning," *arXiv preprint arXiv:1912.11206*, 2019.
- [8] N. Lambert, B. Amos, O. Yadan, and R. Calandra, "Objective mismatch in model-based reinforcement learning," *Learning for Dynamics and Control (LADC)*, pp. 761–770, 2020.
- [9] D. Nguyen-Tuong and J. Peters, "Model learning for robot control: a survey," *Cognitive processing*, vol. 12, no. 4, pp. 319–340, 2011.
- [10] D. H. Shim, H. J. Kim, and S. Sastry, "Decentralized nonlinear model predictive control of multiple flying robots," in *IEEE International Conference on Decision and Control*, vol. 4, 2003, pp. 3621–3626.
- [11] P.-B. Wieber, "Trajectory free linear model predictive control for stable walking in the presence of strong perturbations," in *IEEE-RAS International Conference on Humanoid Robots*, 2006, pp. 137–142.
- [12] G. Klančar and I. Škrjanc, "Tracking-error model-based predictive control for mobile robots in real time," *Robotics and autonomous systems*, vol. 55, no. 6, pp. 460–469, 2007.
- [13] M. P. Deisenroth and C. E. Rasmussen, "PILCO: A Model-Based and Data-Efficient Approach to Policy Search," in *International Conference on Machine Learning*, 2011, pp. 465–472.
- [14] M. Janner, J. Fu, M. Zhang, and S. Levine, "When to trust your model: Model-based policy optimization," in *Neural Information Processing Systems*, 2019, pp. 12498–12509.
- [15] A. Nagabandi, G. Yang, T. Asmar, R. Pandya, G. Kahn, S. Levine, and R. S. Fearing, "Learning image-conditioned dynamics models for control of underactuated legged millirobots," in *IEEE International Conference on Intelligent Robots and Systems*, 2018, pp. 4606–4613.
- [16] N. O. Lambert, D. S. Drew, J. Yaconelli, S. Levine, R. Calandra, and K. S. Pister, "Low-level control of a quadrotor with deep model-based reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 4224–4230, 2019.
- [17] L. Ljung, "System identification," *Wiley encyclopedia of electrical and electronics engineering*, pp. 1–19, 1999.
- [18] A. Venkatraman, M. Hebert, and J. A. Bagnell, "Improving multi-step prediction of learned time series models," in *AAAI Conference on Artificial Intelligence*, 2015.
- [19] E. J. Heravi and S. Khanmohammadi, "Long term trajectory prediction of moving objects using gaussian process," in *IEEE International Conference on Robot, Vision and Signal Processing*, 2011.
- [20] K. Nar, Y. Xue, and A. M. Dai, "Learning unstable dynamical systems with time-weighted logarithmic loss," *arXiv preprint arXiv:2007.05189*, 2020.
- [21] N. R. Ke, A. Singh, A. Touati, A. Goyal, Y. Bengio, D. Parikh, and D. Batra, "Learning dynamics model in reinforcement learning by incorporating the long term future," *arXiv preprint arXiv:1903.01599*, 2019.
- [22] A. Doerr, C. Daniel, D. Nguyen-Tuong, A. Marco, S. Schaal, T. Marc, and S. Trimpe, "Optimizing long-term predictions for model-based policy search," in *Conference on Robot Learning*, 2017, pp. 227–238.
- [23] L. Hewing, E. Arcari, L. P. Fröhlich, and M. N. Zeilinger, "On simulation and trajectory prediction with gaussian process dynamics," in *Learning for Dynamics and Control*, 2020, pp. 424–434.
- [24] A. Wilson, A. Fern, and P. Tadepalli, "Using trajectory data to improve bayesian optimization for reinforcement learning," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 253–282, 2014.
- [25] R. T. Chen, Y. Rubanova, J. Bettencourt, and D. K. Duvenaud, "Neural ordinary differential equations," in *Neural Information Processing Systems*, 2018, pp. 6571–6583.
- [26] D. Ha and J. Schmidhuber, "World models," *arXiv preprint arXiv:1803.10122*, 2018.
- [27] R. Bellman, "A Markovian decision process," *Journal of mathematics and mechanics*, pp. 679–684, 1957.
- [28] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *International Conference on Learning Representations*, 12 2014.
- [29] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Neural Information Processing Systems*, 2014, pp. 3104–3112.
- [30] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *IEEE International Conference on Intelligent Robots and Systems*, 2012, pp. 5026–5033.
- [31] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," *arXiv preprint arXiv:1606.01540*, 2016.
- [32] P. Giernacki, M. Skwarczyński, W. Witwicki, P. Wroński, and P. Kozierski, "Crazyflie 2.0 quadrotor as a platform for research and education in robotics and control engineering," in *IEEE International Conference on Methods and Models in Automation and Robotics (MMAR)*, 2017, pp. 37–42.
- [33] R. Mahony, V. Kumar, and P. Corke, "Multirotor aerial vehicles: Modeling, estimation, and control of quadrotor," *IEEE Robotics and Automation magazine*, vol. 19, no. 3, pp. 20–32, 2012.
- [34] A. Marco, P. Hennig, J. Bohg, S. Schaal, and S. Trimpe, "Automatic LQR tuning based on gaussian process global optimization," in *International Conference on Robotics and Automation*, 2016, pp. 270–277.
- [35] R. Calandra, A. Seyfarth, J. Peters, and M. P. Deisenroth, "Bayesian optimization for learning gaits under uncertainty," *Annals of Mathematics and Artificial Intelligence*, vol. 76, no. 1-2, pp. 5–23, 2016.