

Learning Reactive and Predictive Differentiable Controllers for Switching Linear Dynamical Models

Saumya Saxena¹, Alex LaGrassa¹ and Oliver Kroemer¹

Abstract—Humans leverage the dynamics of the environment and their own bodies to accomplish challenging tasks such as grasping an object while walking past it or pushing off a wall to turn a corner. Such tasks often involve switching dynamics as the robot makes and breaks contact. Learning these dynamics is a challenging problem and prone to model inaccuracies, especially near contact regions. In this work, we present a framework for learning composite dynamical behaviors from expert demonstrations. We learn a switching linear dynamical model with contacts encoded in switching conditions as a close approximation of our system dynamics. We then use discrete-time LQR as the differentiable policy class for data-efficient learning of control to develop a control strategy that operates over multiple dynamical modes and takes into account discontinuities due to contact. In addition to *predicting* interactions with the environment, our policy effectively *reacts* to inaccurate predictions such as unanticipated contacts. Through simulation and real world experiments, we demonstrate generalization of learned behaviors to different scenarios and robustness to model inaccuracies during execution.

I. INTRODUCTION

Everyday human tasks involve intricate interactions with the environment, manifesting in changing dynamics and impacts. For example, to pick up a bottle in one smooth motion while walking past it, a human first slows down, comes in contact, and then applies an appropriate force to pick it up without toppling it over. Modeling the dynamics for such tasks is a challenging problem. Even small inaccuracies in the learned model, especially near the contact regions, can cause catastrophic failures and unstable behaviour. One approach to prevent such problems is to not rely on the learned models entirely [1], [2], rather use them as close approximations of the system’s behaviour, and at the same time design/learn a controller that is both robust to the model inaccuracies and stable under execution noise.

In this work, we develop a control strategy that operates over multiple dynamical modes, takes into account discontinuities due to contact, and is robust to inaccuracies in the learned dynamics model. Our approach focuses on learning the dynamic behavior of an expert and then generalizing that behavior to unseen scenarios.

Systems for which the dynamics change discontinuously, usually after contact is made or broken, can be modeled as *switching linear dynamical systems* (SLDS) [3], [4], [5],

[6], [7], [8]. The set of discrete dynamics in a SLDS are governed by hidden states called *modes*. We model impacts as conditions governing the switch between these modes, which incorporates discontinuities into our dynamics model. We use system identification to learn the parameters of the SLDS including the impact dynamics and mode prediction function.

Taking inspiration from recent approaches that add prior structure to policy architecture [9], [10], [11] for data-efficient learning of control, we use discrete-time LQR as our differentiable policy class. In addition to being sample efficient, this policy class provides a feedback control scheme that captures expert behavior in the learned cost/value functions which can be used to generalize to other scenarios. For our complex task modeled as a SLDS, we use a single LQR policy and parameterize it using a different cost function and goal condition for each mode. Using cost matrices learned from a single scenario, we can generalize to different goal conditions while achieving the same overall behavior.

Forward propagating the control using the learned model allows us to predict mode switches and contacts by extension. To make the controller robust to inaccurate mode predictions, we combine our method with a reactive control scheme [12], to switch modes upon *observing* inaccurate predictions of state transitions. We use model predictive control during execution to replan using the observed modes.

The contributions of this paper are twofold: 1) a method for modeling and learning switching contact-based linear dynamics and using it for efficiently learning a differentiable feedback controller that operates over multiple dynamical modes, takes into account discontinuities due to contact, and generalizes it to unseen scenarios, 2) a control scheme that is robust to inaccuracies in the learned dynamics model and during execution can predict, as well as react to, unanticipated dynamic contact events. An overview of the method and results can be found in the supplementary video [\(link\)](#).

II. RELATED WORK

Contacts and changing dynamics are ubiquitous in manipulation [13] and collaborative tasks. Model-based control approaches for learning contact-rich manipulation have been proposed [14], [15], but their asymptotic performance suffers due to model-bias. To alleviate this problem, other approaches combine model-based and model-free methods [1], [16]. Our approach relies on using a robust control scheme [17] to deal with inaccuracies in the learned dynamics model.

Modeling tasks with non-stationary dynamics as switching linear dynamical systems (SLDS) is a well known problem.

*This work was in part supported by the National Science Foundation under Grant No. CMMI-1925130 and IIS-1956163. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.

¹Robotics Institute, School of Computer Science, Carnegie Mellon University, Pittsburgh PA 15123, USA {saumyas, alagrass, okroemer}@andrew.cmu.edu

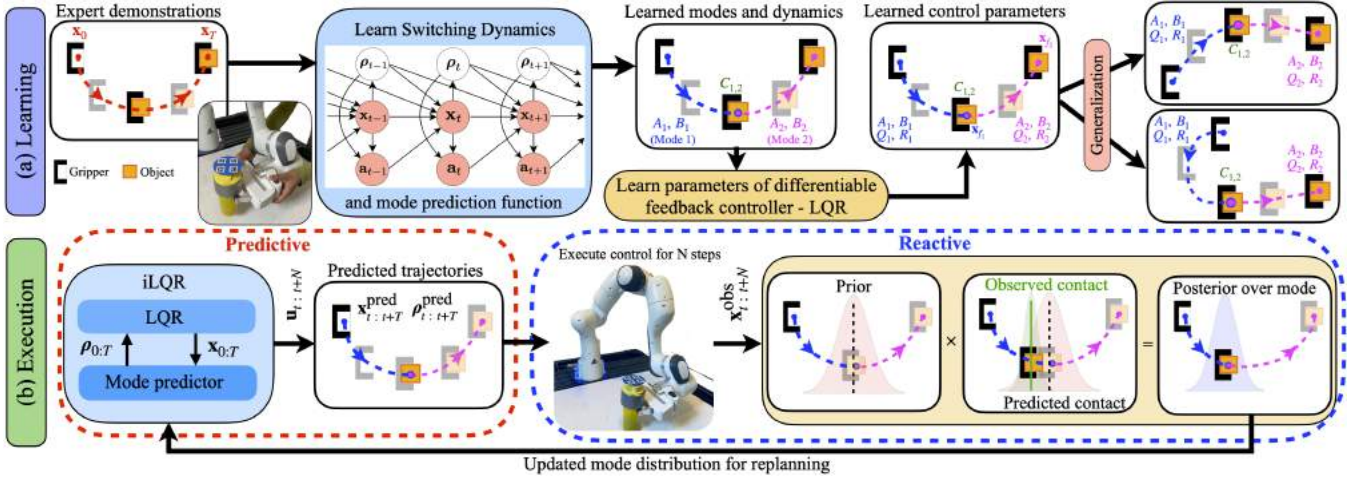


Fig. 1. Training and execution pipeline (a) Given a set of expert demonstrations, learn the parameters of a switching linear dynamical system (transition dynamics, contact dynamics and mode prediction function) using system identification. Using the learned dynamics parameters and the expert demonstrations, learn a differentiable feedback controller (LQR), parameterized using a different cost function and goal condition for each mode, using imitation learning. Learning the cost matrices allows us to capture the behavior of the expert in each mode, thus enabling the learned controller to generalize to unseen scenarios. (b) During execution, we use iLQR to iteratively evaluate the modes along a trajectory, and use these modes and hence the corresponding dynamics, impact and cost parameters to solve the LQR problem. To make the controller robust to inaccuracies in the learned model, we use MPC to combine our predictive controller with a reactive control scheme. The predictive part adapts the control in anticipation of mode switches, and the reactive part observes unanticipated mode transitions, updates the current mode and replans.

There are approaches that do [7], [8] and do not [3], [4], [5], [6] model discontinuous mode transitions. Some other approaches for modelling contacts and learning contact based dynamics include [18], [19], [20], [21]. Our approach for modeling contact as mode switches compares best with [22], where the mapping from pre-impact velocity to post-impact velocity is learned. Methods such as [3], [23] learn controllers for each mode separately, and switch between controllers using the learned predictive guard conditions. Such methods are susceptible to failures due to inaccuracies in these learned conditions. Methods such as eMOSAIC [12], on the other hand, do not learn a predictive model for mode switching, but switch modes and control reactively upon observed changes in the effects of actions. In our method, we combine the advantages of both the predictive and reactive control schemes to develop a robust controller that acts over multiple discrete modes and contact.

Model-free approaches [24], [25], [26] obviate the need for learning the dynamics, but face challenges in terms of sample complexity. Recent approaches to alleviate these challenges propose adding prior structure to the policy architecture [9], [10], [11]. We take inspiration from this work and use discrete-time LQR as our differentiable policy class.

III. METHOD

In this section we first elaborate on our method for modeling and learning parameters of the switching dynamical system using system identification. We then model the controller and learn the control parameters using imitation learning. Finally we introduce our robust control scheme for executing the learned policy.

A. Modeling the Switching Linear Dynamical System

Graphical model for our switching dynamical system is shown in the second panel in Fig. 1(a). State of the system at time t is given by $\mathbf{x}_t = [\mathbf{q}_t, \dot{\mathbf{q}}_t]$ where $\mathbf{x}_t \in \mathbb{R}^n$, $\mathbf{q}_t, \dot{\mathbf{q}}_t$ are

configuration and velocities respectively. State is composed of state of the robot and states of objects of interest in the environment. At every state \mathbf{x}_t the robot takes an action $\mathbf{u}_t \in \mathbb{R}^m$ and transitions to state \mathbf{x}_{t+1} . The transition dynamics and the control parameters at every time step are dependent on the discrete mode ρ_t , which is a hidden variable.

1) *Mode Prediction:* We assume that the current mode ρ_t is only dependent upon the current state \mathbf{x}_t . We call the function for mode prediction a *classifier* $\rho_t = M(\mathbf{x}_t)$. The classifier is modeled as a fully-connected neural network that takes as input current state \mathbf{x}_t and outputs a one-hot vector over modes. For a picking task the state \mathbf{x}_t will include the position and orientation of the object we want to pick up. The classifier can learn to predict modes that depend upon more general conditions such as relative position between the gripper and the object, as used in our model. We always set the number of modes in the model to be greater or equal to the actual number of modes required by the task. The number of learned modes can be less than or equal to the maximum value specified.

2) *Transition and Interaction dynamics:* The transition dynamics are modeled as follows:

$$\mathbf{x}_{t+1} = \mathbf{A}_{\rho_t} \mathbf{C}_{\rho_t, \rho_{t-1}} \mathbf{x}_t + \mathbf{B}_{\rho_t} \mathbf{u}_t \quad (1)$$

where $\mathbf{A}_{\rho_t} \in \mathbb{R}^{n \times n}$ and $\mathbf{B}_{\rho_t} \in \mathbb{R}^{n \times m}$ correspond to the discrete mode ρ_t at time t . Thus, for a set of N modes, there will be N such transition matrices that we need to learn. The matrix $\mathbf{C}_{\rho_t, \rho_{t-1}} \in \mathbb{R}^{n \times n}$ represents the impact or interaction dynamics which depends upon the mode switch between the current and previous time step. For example, in case of a pushing task with two modes, the first mode is when the robot is in free motion and the second mode is when it is in contact with the light object and is pushing it. An impact occurs between the robot and the object at the instant when contact is made, that is, when the mode switches from

mode 1 to 2. Upon impact, the velocity of a system changes instantaneously while the position remains the same, which can be modeled as momentum transfer between the robot and the object. This is the source of discontinuity in the dynamics which we model here explicitly.

We take inspiration from approaches that model contact as a mapping from pre-impact to post-impact velocity [22]. Consider a simple system where the gripper and object states are given as $\mathbf{x}_t^g = [\mathbf{q}_t^g, \dot{\mathbf{q}}_t^g]$ and $\mathbf{x}_t^o = [\mathbf{q}_t^o, \dot{\mathbf{q}}_t^o]$, respectively. The impact dynamics for this system can be written as

$$\begin{bmatrix} \mathbf{q}_{t+}^g \\ \dot{\mathbf{q}}_{t+}^g \\ \mathbf{q}_{t+}^o \\ \dot{\mathbf{q}}_{t+}^o \end{bmatrix} = \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \frac{m_g - em_o}{m_g + m_o} & 0 & \frac{(1+e)m_o}{m_g + m_o} \\ 0 & 0 & 1 & 0 \\ 0 & \frac{(1+e)m_g}{m_g + m_o} & 0 & \frac{m_o - em_g}{m_g + m_o} \end{bmatrix}}_{\mathbf{C}} \begin{bmatrix} \mathbf{q}_{t-}^g \\ \dot{\mathbf{q}}_{t-}^g \\ \mathbf{q}_{t-}^o \\ \dot{\mathbf{q}}_{t-}^o \end{bmatrix},$$

where e is the coefficient of restitution between the gripper and the object, m_g and m_o are the mass matrices for the gripper and the object respectively, and $\mathbf{q}_{t-}, \dot{\mathbf{q}}_{t-}$ and $\mathbf{q}_{t+}, \dot{\mathbf{q}}_{t+}$ are the position and velocity before and after impact respectively. We observe that some properties of this impact/interaction matrix \mathbf{C} make it easier to learn from demonstration and apply it in a wide range of applicable scenarios. First, \mathbf{C} is only dependent on the system properties and not on the velocity or position. Thus, we can parameterize and learn it using expert demonstrations and then use it in different scenarios. Second, we need only learn the parameters corresponding to changes in velocity because the position remains the same during impact. Finally, even though we model a different impact matrix for each possible mode switch, we end up learning impact matrices only for the mode switches that are observed during demonstrations, pruning the number of impact matrices that need to be learned. To ensure that the impact only occurs upon a mode *switch*, we write the impact dynamics as $\mathbf{C}_{\rho_{t-1}, \rho_t} = \rho_{t-1}^\top \rho_t \mathbf{I} + (1 - \rho_{t-1}^\top \rho_t) \mathbf{C}$ because when $\rho_{t-1} = \rho_t$, $\mathbf{C}_{\rho_{t-1}, \rho_t} = \mathbf{I}$ and when $\rho_{t-1} \neq \rho_t$, $\mathbf{C}_{\rho_{t-1}, \rho_t} = \mathbf{C}$ where \mathbf{C} is learned from demonstrations.

B. Learning the Switching Dynamics Model

Having defined the structure of the dynamical system, we move on to learning the model parameters from expert demonstrations using system identification. The expert demonstrations are given as a set of trajectories $\{\mathbf{x}_t^*, \mathbf{u}_t^*, \mathbf{x}_{t+1}^*\}_{t=\{0, \dots, T-1\}}$ where T is the total length of the trajectory. The system identification cost is written as:

$$\rho_t = M(\mathbf{x}_t^*), \quad \mathbf{x}_{t+1}^{\text{pred}} = \mathbf{A}_{\rho_t} \mathbf{C}_{\rho_t, \rho_{t-1}} \mathbf{x}_t^* + \mathbf{B}_{\rho_t} \mathbf{u}_t^* \\ \mathcal{L}_{\text{SysID}} = \|\mathbf{x}_{1:T}^* - \mathbf{x}_{1:T}^{\text{pred}}\|_2.$$

Taking gradient with respect to the above cost we together learn $\mathbf{A}_{\rho_t}, \mathbf{B}_{\rho_t}, \mathbf{C}_{\rho_t, \rho_{t-1}}$, and the classifier $M(\cdot)$.

C. Modeling the Differentiable Policy

Having learned the system dynamic parameters, we shift our focus to modeling a controller that can perform the task demonstrated by the expert. More importantly, the

controller should 1) robustly perform the task, 2) match the behavior of the expert through the learned cost function, such as slowing down before grasping, 3) generalize to other task instances that use the same objects. We design our controller as a discrete-time linear quadratic regulator (LQR), a differentiable feedback controller that can be represented as the solution to the following optimization problem:

$$\begin{aligned} \min_{\mathbf{u}_t} \quad & \sum_{t=0}^T ((\mathbf{x}_t - \mathbf{x}_{f_{\rho_t}})^\top \mathbf{Q}_{\rho_t} (\mathbf{x}_t - \mathbf{x}_{f_{\rho_t}}) + \mathbf{u}_t^\top \mathbf{R}_{\rho_t} \mathbf{u}_t) \\ \text{s.t.} \quad & \mathbf{x}_{t+1} = \mathbf{A}_{\rho_t} \mathbf{C}_{\rho_t, \rho_{t-1}} \mathbf{x}_t + \mathbf{B}_{\rho_t} \mathbf{u}_t, \quad i = 1, \dots, T \\ & \mathbf{x}_0 = \mathbf{x}_{\text{init}} \end{aligned} \quad (2)$$

where $\mathbf{x}_{f_{\rho_t}}, \mathbf{Q}_{\rho_t}, \mathbf{R}_{\rho_t}$ correspond to the goal configuration and cost matrices for mode ρ_t . This formulation allows us to capture the behavior of the system in each mode using the cost matrices \mathbf{Q} and \mathbf{R} for a corresponding goal location \mathbf{x}_f . It is important to note that the behavior over the entire task is governed by the sequence of behaviors of the learned modes. Thus, once learned from expert demonstrations, as illustrated in Section III-B, the *sequence* of modes should be kept fixed to successfully complete the task.

Using LQR feedback controller enables optimal responses to anticipated changes in dynamics and impact conditions in the trajectory. By learning cost matrices, the agent is essentially learning a value function or reward structure around each goal configuration. Using learned cost matrices allows the controller to generalize to different sub-goal configurations along the trajectory. Keeping the sequence of modes, and thus the *behavior* along the trajectory, fixed, we can shift around the learned sub-goal locations (\mathbf{x}_f) to new desired locations and still attain the same overall behavior.

D. Learning the Control Parameters

As illustrated in Section III-B, for each expert demonstration, we have already learned the mode prediction $\rho_t = M(\mathbf{x}_t^*)$ and the corresponding dynamic parameters $\mathbf{A}_{\rho_t}, \mathbf{B}_{\rho_t}, \mathbf{C}_{\rho_t, \rho_{t-1}}$ for each time step. Now, we use imitation learning to learn the control parameters $\mathbf{x}_{f_{\rho}}, \mathbf{Q}_{\rho}$ and \mathbf{R}_{ρ} . For each expert initial condition $\{\mathbf{x}_0^*\}_{t=\{0, \dots, T\}}$, using the corresponding learned dynamic parameters and modes, we solve the optimization problem in (2) using discrete-time Riccati equations [27]. Let the solution to the LQR problem for each initial state \mathbf{x}_0^* be given as

$$\mathbf{x}_{0:T}^{\text{LQR}}, \mathbf{u}_{0:T-1}^{\text{LQR}} = \text{LQR}(\mathbf{x}_0^*, \mathbf{A}_{\rho_t}, \mathbf{B}_{\rho_t}, \mathbf{C}_{\rho_t, \rho_{t-1}}, \mathbf{Q}_{\rho_t}, \mathbf{R}_{\rho_t}, \mathbf{x}_{f_{\rho_t}}).$$

The imitation cost can then be written as

$$\mathcal{L}_{\text{Imitation}} = \|\mathbf{x}_{0:T}^* - \mathbf{x}_{0:T}^{\text{LQR}}\|_2 + \|\mathbf{u}_{0:T-1}^* - \mathbf{u}_{0:T-1}^{\text{LQR}}\|_2.$$

The differentiability of the LQR controller allows us to learn the goal configurations $\mathbf{x}_{f_{\rho}}$ and cost functions $\mathbf{Q}_{\rho}, \mathbf{R}_{\rho}$ corresponding to each mode by minimizing the above cost.

E. Execution

We have shown how we learn the dynamics model and controller from expert demonstrations. Now we move on to

how we can solve similar tasks with new initial configurations of the robot and objects in the environment using these learned parameters. We also show how the controller accounts for inaccuracies in the learned dynamics model.

Algorithm 1 Iterative LQR

(i) represents the iteration number

T is the length of the trajectory

$j, k \in \{1, \dots, N\}$ where N is the number of modes

Given: Initial state \mathbf{x}_{init} . Initialize modes $\rho_{0:T}$

Parameters: $M(\cdot), \mathbf{A}_{\rho_j}, \mathbf{B}_{\rho_j}, \mathbf{C}_{\rho_j}, \mathbf{C}_{\rho_j, \rho_k}, \mathbf{x}_{f_{\rho_j}}, \mathbf{Q}_{\rho_j}, \mathbf{R}_{\rho_j}$

procedure iLQR(Parameters)

while not converged **do**

$\mathbf{x}_{0:T}^{(i)}, \mathbf{K}_{0:T-1} \leftarrow \text{LQR}(\mathbf{A}_{\rho_t^{(i)}}, \mathbf{B}_{\rho_t^{(i)}}, \mathbf{C}_{\rho_t^{(i)}}, \mathbf{C}_{\rho_t^{(i)}, \rho_{t-1}^{(i)}}, \mathbf{x}_{f_{\rho_t^{(i)}}}, \mathbf{Q}_{\rho_t^{(i)}}, \mathbf{R}_{\rho_t^{(i)}})$

$\rho_t^{(i+1)} = M(\mathbf{x}_t^{(i)})$ for all t

 converged $\leftarrow ||\mathbf{x}_{0:T}^{(i+1)} - \mathbf{x}_{0:T}^{(i)}||_2 + ||\mathbf{u}_{0:T-1}^{(i+1)} - \mathbf{u}_{0:T-1}^{(i)}||_2$ is small

end while

end procedure

1) *Iterative Linear Quadratic Regulator:* We formulate the problem of evaluating the control sequence for tasks involving multiple discrete dynamical modes as an iterative LQR problem. The approach is very similar to a standard iLQR problem where, starting with an initial trajectory, we linearize the dynamics around the trajectory and use these dynamics to solve the optimal LQR problem iteratively. In our approach, as illustrated in Algorithm 1, at every iteration, we evaluate the mode for each step of the trajectory and populate the dynamics, impact and cost matrices accordingly to iteratively solve the LQR problem. It is important to note here that, from the expert demonstrations, we learned the cost matrices that capture the behavior of the system while achieving goal states $\mathbf{x}_{f_{\rho_j}}$. By substituting $\mathbf{x}_{f_{\rho_j}}$ with new desired goal states, we can generalize to different initial and goal configurations for a task with the same sequence of modes. A high-level planner can be used to rearrange the modes to generalize to similar tasks with different mode sequences, but this has been left for future work.

Algorithm 2 Model Predictive Control

T_{LQR} is time horizon for execution before replanning

$N_{MPC} = T/T_{LQR}$

T is the length of the trajectory

t is the current time step

$t \leftarrow 0$

for n in N_{mpc} **do**

$\mathbf{K}_{t:T} \leftarrow \text{iLQR}(\mathbf{x}_t, \mathbf{A}_{\rho}, \mathbf{B}_{\rho}, \mathbf{C}_{\rho}, \mathbf{C}_{\rho, \rho}, \mathbf{R}_{\rho}, \mathbf{Q}_{\rho}, \mathbf{x}_{f_{\rho}}, T-t)$

for i in T_{LQR} **do**

$\rho_t^{post} \leftarrow \text{Posterior}(\mathbf{x}_t, \mathbf{x}_{t-1}, \mathbf{a}_{t-1})$

$\mathbf{a}_t \leftarrow \mathbf{K}_t(\mathbf{x}_t - \mathbf{x}_{f_{\rho_t}}^{post})$

$\mathbf{x}_{t+1} \leftarrow \text{step}(\mathbf{a}_t)$

$t \leftarrow t_{t+1}$

end for

end for

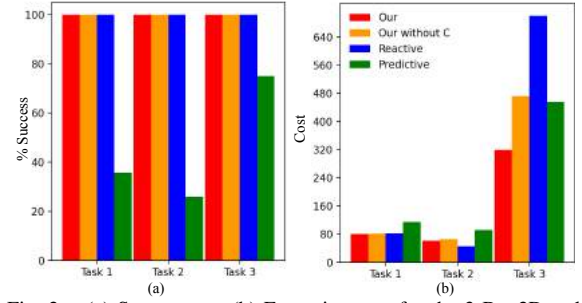


Fig. 2. (a) Success rate (b) Execution cost for the 3 Box2D tasks

Table I
RMSE of learned masses and learned goal configurations.

	Masses (kg)	Goal configurations (m)
Task 1	0.029	0.095
Task 2	0.034	0.088
Task 3	0.009	0.026

2) *Model Predictive Control:* Since LQR is a feedback controller, it is robust to small disturbances and noise that are encountered during execution. However, our controller should be robust to inaccuracies in our learned model, especially the learned mode predictor $M(\mathbf{x}_t)$. Consider a scenario where for a grasping task, the classifier learns that contact occurs when the measured relative position between the gripper and the object is less than some threshold value. During execution, contact could potentially occur a little before, or a little after, this threshold is crossed. If the threshold is a little too large and the next mode requires the robot to move away from the object location, the robot might start moving away before actually grasping the object. To make the system robust to such scenarios, we use model predictive control (MPC), Algorithm 2, to *replan* using observations from the environment. We evaluate a posterior over the mode estimate using the current observation as

$$p(\rho_t | \mathbf{x}_t, \mathbf{u}_{t-1}, \mathbf{x}_{t-1}) = z * p(\rho_t | \mathbf{x}_t) p(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{u}_{t-1}, \rho_t) \quad (3)$$

where $p(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{u}_{t-1}, \rho_t)$ gives the probability of seeing the observed transition given a certain mode ρ_t , $p(\rho_t | \mathbf{x}_t)$ is the predicted mode from the learned classifier, and z is the normalization constant. This approach to evaluating the mode using observed state transitions is similar to the eMOSAIC approach [12]. The above posterior can also be calculated using measured contacts if the robot is equipped with force/contact sensing. Using this posterior estimate of the mode, we replan as given in Algorithm 2. Adding this framework to our approach, we enable our controller to be predictive of mode transitions (using the classifier) as well as reactive (using measurements to evaluate posterior) to unanticipated mode changes and impacts conditions making the controller robust to learned model inaccuracies.

IV. EXPERIMENTS

The proposed method was evaluated in simulation using the Box2D environment and in the real-world on the 7DOF Franka-Emika Panda robot arm.

A. Simulation Experiments

To evaluate our controller, which is both reactive and predictive and takes into account discontinuities due to

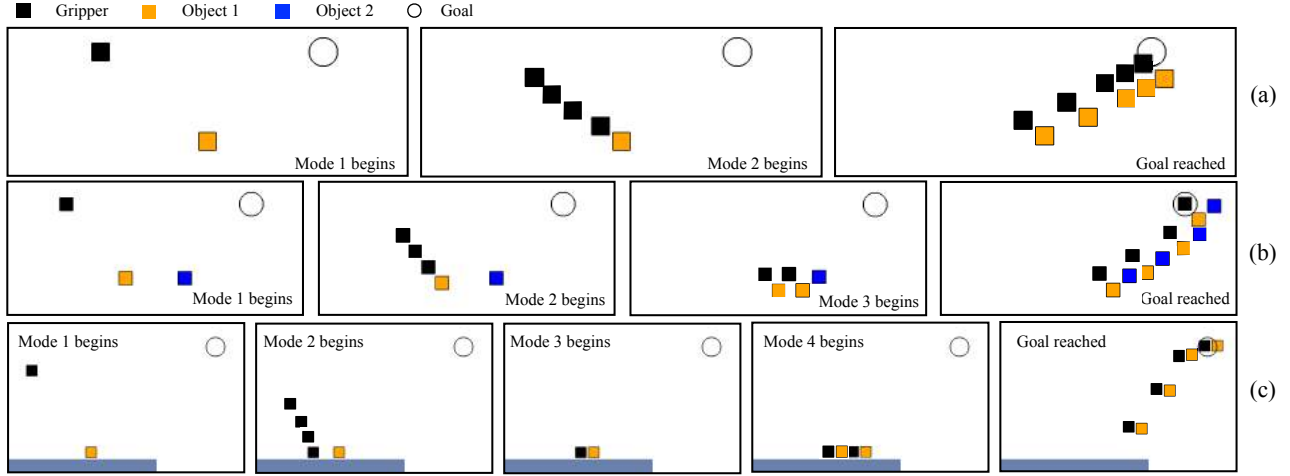


Fig. 3. Learned behaviors for (a) 1-object pickup task, (b) 2-object pickup task, and (c) pickup task requiring sliding along the table to align with the object before picking it up. A 2D gripper (black block) picks up objects (orange and blue blocks) and takes it to goal region (black circle). Motion of gripper and objects is depicted by overlaid transparent frames.

contact, we compare our results to 1) a purely reactive eMOSAIC controller [12], 2) a purely predictive version of our controller, i.e. without using the posterior over the mode to replan using MPC, and 3) a predictive and reactive controller that doesn't take into account discontinuities due to contact. Experiments were designed to test the robustness of these controllers to process noise, model inaccuracies and variations in shape geometry and initial configurations not seen during training. We use two metrics to compare results: 1) success rate - defined by whether the robot is able to pick up the object(s) of interest and take them to a goal region in a specified time 2) cost to perform the task, defined as the LQR cost: $\sum_{t=0}^T ((\mathbf{x}_t - \mathbf{x}_{f_{\rho_t}})^T \mathbf{Q}_{\rho_t} (\mathbf{x}_t - \mathbf{x}_{f_{\rho_t}}) + \mathbf{u}_t^T \mathbf{R}_{\rho_t} \mathbf{u}_t)$. We perform three tasks in simulation: 1) pick up one object and deliver to goal location 2) pick up two objects and deliver to goal location 3) pick up task while interacting with the ground to align with the object.

Expert demonstrations are collected by sampling the initial configuration for the gripper and fixed initial configurations for the objects. We use ground truth dynamic properties of the objects and manually defined cost functions to collect expert demonstrations for each task in simulation. By hand-tuning the cost matrices, we can control how much the gripper slowed down before picking up each object. For each of the tasks, the classifier and dynamics model are learned first using system identification as described in Section III-B. Then, the cost matrices and goal conditions are learned as described in Section III-D. For the eMOSAIC controller, these learned cost matrices and goal conditions are used to design independent infinite-horizon LQR controllers, one for each mode. The root mean squared error (RMSE) between the learned and ground truth values for the masses and goal configurations for each mode are summarized in Table II. Additional details concerning training and experiments are provided in appendix (link).

1) *Task 1: 1-object pickup*: We consider the task of picking up an object using a 2D gripper and taking it to a goal region, specified as the black circle shown in Fig. 3a. The classifier learns to predict two modes. Mode 1 is

active when the gripper is away from the object and mode 2 is active when the relative position between the gripper and the object is smaller than a learned threshold value. For this task, we observe that the classifier learns to predict contact a little after the actual impact occurs signifying imperfectness in the learned model. The masses for the gripper and object are set as 1kg and 9kg respectively in simulation. We perform multiple trials with varying initial configurations and shape parameters. As shown in Fig. 2, our controller with and without impact dynamics ($\mathbf{C}_{\rho_t, \rho_{t-1}}$) and eMOSAIC generalize well to new scenarios. Our controller is robust to local inaccuracies and eMOSAIC is unaffected by them as it is purely reactive. Both of these methods are successful in taking the object to the goal location in all the trials. On the other hand, performance of the predictive controller suffers because of inaccuracy in the learned classifier. The predictive controller fails to identify a contact even after the contact has actually occurred and does not switch to the next mode. We observe that the trajectories generated by eMOSAIC are non-smooth and the impact velocities are much higher than the expert's, leading to higher cost. In contrast, the resulting cost of our controller is lower than the other controllers because of its capability to predict the next mode, curving the trajectory of the gripper in the direction of the next goal early on. The cost of our controller without $\mathbf{C}_{\rho_t, \rho_{t-1}}$ is also slightly higher as it doesn't consider the discontinuous jump in velocities of the object and gripper upon contact.

2) *Task 2: 2-object pickup*: We consider the task of picking up two objects using a 2D gripper and moving them to a goal region, as shown in Fig. 3b. The model learns to identify three modes: Mode 1) relative position between the gripper and both objects is greater than a threshold value, Mode 2) relative position between the gripper and object 1 (orange) is smaller than a threshold value whereas between gripper and object 2 (blue) is larger than a threshold value, and Mode 3) relative position between the gripper and both objects is smaller than a threshold value. The masses of the gripper and objects in simulation were set as 1kg each. For this task, we observe that the learned classifier predicts



Fig. 4. Sequence of images showing the Franka robot arm performing a dynamic pickup task using our learned reactive and predictive controller. The controller generalizes to previously unseen locations of the object and performs the task with 80% success rate.

contacts a little earlier than when contact actually happens. As in Task 1, we observe that our controller with and without $\mathbf{C}_{\rho_t, \rho_{t-1}}$ is robust to these local inaccuracies, eMOSAIC is unaffected by it and both generalize well to new scenarios. As shown in Fig. 2 both these controllers were able to successfully perform the task in all trials. The success rate for the predictive controller suffers because in a large number of trials the gripper moves past the object without actually grasping it because the classifier wrongly predicted that grasp has happened. eMOSAIC performs better than our controller for this task because it is purely predictive. The inaccurate classifier leads our method to predict suboptimal trajectories in the short term, which leads to higher cost.

3) *Task 3: Pickup task with ground interactions:* This experiment evaluates the capability of our method to capture *intricate behaviors* from expert demonstrations. These demonstrations were designed to depict behavior that exploits contact to reduce uncertainty. We consider the task of picking up an object placed on a table. Fig. 3 demonstrates how the gripper first comes in contact with the table and moves along it before coming in contact with the object and pushing it along the table before picking it up. As before, we perform experiments by varying initial configurations and block sizes. We observe that the success rate for the predictive controller improves as compared to other tasks because alignment with the object before picking it up reduces uncertainty in relative positions. Our controller with and without $\mathbf{C}_{\rho_t, \rho_{t-1}}$ and eMOSAIC are successful in completing the task in every trial but the cost increases significantly if $\mathbf{C}_{\rho_t, \rho_{t-1}}$ is not considered. This is because the controller is unable to predict that the velocity of the gripper in the direction of contact goes to zero upon contact with the ground and the subsequent motion is along the ground. This leads to suboptimal predictions and thus higher cost. The cost for eMOSAIC is also high in this task. This is because the gripper and object overshoot the table due to the lack of a predictive component preparing the gripper to move up towards the goal early on, requiring a large force to correct the motion to eventually move toward the goal.

B. Robot Experiments

We consider the task of dynamically picking up an object and taking it to a goal location as shown in Fig. 4. Ten expert demonstrations were collected using human-guided kinesthetic teaching [28] with fixed object and goal locations. In the human demonstration, the robot slows down before contact, without stopping, to prevent the object from toppling. The location of the object was tracked using April tags and a Kinect depth sensor [29]. The robot dynamics were linearized using Cartesian space feedback linearization.

Due to a lack of reliable contact sensing, we only rely on sensing of the object location as feedback to compare the performance of our controller with the reactive controller.

Using the method described in Section III-B, the classifier learns to discretize the task into two modes. The goal condition for the first mode (coming in contact with object) was learned with RMSE of 0.015m, and the goal condition for the second mode (final configuration) was learned with a RMSE of 0.079m. To compare our controller to the reactive eMOSAIC controller we perform 10 trials with varying initial locations of object. Using our controller, the robot successfully picks up the object and takes it to the goal location 8/10 times. Average cost over successful trials with our controller is 1.1789. Using the purely reactive eMOSAIC controller, the robot successfully picks up the object and takes it to goal location 6/10 times. The average cost over the successful trials is 1.272. The sharp movements of the reactive controller right after coming in contact with the object is the main reason behind the failed attempts. Unlike in Box2D experiments, where abrupt motion in any direction can occur, such behavior does not perform well in real-world scenarios and can lead to failed grasps. As compared to eMOSAIC, our controller executes smooth trajectories before and after impact as it operates over all modes.

V. CONCLUSIONS

We presented a framework for developing a robust differentiable feedback controller for learned switching linear dynamical systems that operates over multiple dynamical modes, predicts and reacts to discontinuous mode switches and generalizes to unseen scenarios. Expert demonstrations inform mode predictions, contact dynamics, linear dynamical system parameters, cost functions and goal conditions. Experimental results show a significant improvement over eMOSAIC, a purely reactive controller, in multi-object pick-up, sliding and grasping, and real-robot grasping. We find that replanning using MPC and posterior modes improves robustness to model inaccuracies. In future work we will explore planning over modes to generalize to tasks that might require different mode sequencing. We plan to extend the capabilities of our controller to non-linear dynamical systems by learning locally linear dynamical models.

ACKNOWLEDGMENT

The authors would like to thank Kevin Zhang, Blake Buchanan and Jacky Liang for their support during this work. The authors also would like to thank Prof. Zico Kolter for engaging technical discussions on differentiable control algorithms that form the basis of this work.

REFERENCES

- [1] A. Nagabandi, G. Kahn, R. S. Fearing, and S. Levine, "Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 7559–7566.
- [2] M. Janner, J. Fu, M. Zhang, and S. Levine, "When to trust your model: Model-based policy optimization," in *Advances in Neural Information Processing Systems*, 2019, pp. 12519–12530.
- [3] O. Kroemer, C. Daniel, G. Neumann, H. Van Hoof, and J. Peters, "Towards learning hierarchical skills for multi-phase manipulation tasks," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 1503–1510.
- [4] P. Becker-Ehmck, J. Peters, and P. Van Der Smagt, "Switching linear dynamics for variational bayes filtering," *arXiv preprint arXiv:1905.12434*, 2019.
- [5] S. W. Linderman, A. C. Miller, R. P. Adams, D. M. Blei, L. Paninski, and M. J. Johnson, "Recurrent switching linear dynamical systems," *arXiv preprint arXiv:1610.08466*, 2016.
- [6] J. J. Dabrowski, J. P. de Villiers, and C. Beyers, "Naïve bayes switching linear dynamical system: A model for dynamic system modelling, classification, and information fusion," *Information Fusion*, vol. 42, pp. 75–101, 2018.
- [7] M. Toussaint and S. Vijayakumar, "Learning discontinuities with products-of-sigmoids for switching between local models," in *Proceedings of the 22nd international conference on Machine Learning*, 2005, pp. 904–911.
- [8] G. Lee, Z. Marinho, A. M. Johnson, G. J. Gordon, S. S. Srinivasa, and M. T. Mason, "Unsupervised learning for nonlinear piecewise smooth hybrid systems," *arXiv preprint arXiv:1710.00440*, 2017.
- [9] B. Amos, I. Jimenez, J. Sacks, B. Boots, and J. Z. Kolter, "Differentiable mpc for end-to-end planning and control," in *Advances in Neural Information Processing Systems*, 2018, pp. 8289–8300.
- [10] P. Karkus, X. Ma, D. Hsu, L. P. Kaelbling, W. S. Lee, and T. Lozano-Pérez, "Differentiable algorithm networks for composable robot learning," *arXiv preprint arXiv:1905.11602*, 2019.
- [11] M. Toussaint, K. Allen, K. A. Smith, and J. B. Tenenbaum, "Differentiable physics and stable modes for tool-use and manipulation planning," in *Robotics: Science and Systems*, 2018.
- [12] N. Sugimoto, J. Morimoto, S.-H. Hyon, and M. Kawato, "The emosaic model for humanoid robot control," *Neural Networks*, vol. 29, pp. 8–19, 2012.
- [13] O. Kroemer, S. Niekum, and G. Konidaris, "A review of robot learning for manipulation: Challenges," *Representations, and Algorithms*, (2019), 2019.
- [14] S. Gu, T. Lillicrap, I. Sutskever, and S. Levine, "Continuous deep q-learning with model-based acceleration," in *International Conference on Machine Learning*, 2016, pp. 2829–2838.
- [15] N. Mishra, P. Abbeel, and I. Mordatch, "Prediction and control with temporal segment models," *arXiv preprint arXiv:1703.04070*, 2017.
- [16] L. Sergey, N. Wagener, and P. Abbeel, "Learning contact-rich manipulation skills with guided policy search," in *Proceedings of the 2015 IEEE International Conference on Robotics and Automation (ICRA)*, Seattle, WA, USA, 2015, pp. 26–30.
- [17] G. E. Dullerud and F. Paganini, *A course in robust control theory: a convex approach*. Springer Science & Business Media, 2013, vol. 36.
- [18] S. Pfrommer, M. Halm, and M. Posa, "Contactnets: Learning of discontinuous contact dynamics with smooth, implicit representations," *arXiv preprint arXiv:2009.11193*, 2020.
- [19] N. Fazeli, R. Kolbert, R. Tedrake, and A. Rodriguez, "Parameter and contact force estimation of planar rigid-bodies undergoing frictional contact," *The International Journal of Robotics Research*, vol. 36, no. 13-14, pp. 1437–1454, 2017.
- [20] R. Calandra, S. Ivaldi, M. P. Deisenroth, E. Rueckert, and J. Peters, "Learning inverse dynamics models with contacts," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 3186–3191.
- [21] S. A. Khader, H. Yin, P. Falco, and D. Kragic, "Data-efficient model learning and prediction for contact-rich manipulation tasks," *IEEE Robotics and Automation Letters*, vol. 5, no. 3, pp. 4321–4328, 2020.
- [22] N. Fazeli, S. Zapolsky, E. Drumwright, and A. Rodriguez, "Learning data-efficient rigid-body contact models: Case study of planar impact," in *Conference on Robot Learning*, 2017, pp. 388–397.
- [23] S. Manschitz, J. Kober, M. Gienger, and J. Peters, "Learning movement primitive attractor goals and sequential skills from kinesthetic demonstrations," *Robotics and Autonomous Systems*, vol. 74, pp. 97–107, 2015.
- [24] O. M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, *et al.*, "Learning dexterous in-hand manipulation," *The International Journal of Robotics Research*, vol. 39, no. 1, pp. 3–20, 2020.
- [25] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski, *et al.*, "Human-level control through deep reinforcement learning," *nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [26] S. Gu, T. Lillicrap, Z. Ghahramani, R. E. Turner, and S. Levine, "Q-prop: Sample-efficient policy gradient with an off-policy critic," *arXiv preprint arXiv:1611.02247*, 2016.
- [27] A. Bemporad, M. Morari, V. Dua, and E. N. Pistikopoulos, "The explicit linear quadratic regulator for constrained systems," *Automatica*, vol. 38, no. 1, pp. 3–20, 2002.
- [28] K. Zhang, M. Sharma, J. Liang, and O. Kroemer, "A modular robotic arm control stack for research: Franka-interface and frankapy," *arXiv preprint arXiv:2011.02398*, 2020.
- [29] J. Wang and E. Olson, "Apriltag 2: Efficient and robust fiducial detection," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2016, pp. 4193–4198.
- [30] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

VI. APPENDIX

A. Training details

The classifier, for all the experiments, is a single layer neural network with 8 hidden nodes. The output layer has nodes equal to the number of modes and softmax activation. The other hyperparameters used during training (system identification as well as learning control parameters) are as follows:

Table II
Hyperparameters

Optimizer	ADAM [30]
Learning rate	0.01
β_1	0.9
β_2	0.999
ϵ	1e-04
Length of each trajectory	12.5sec
dt	0.05sec
batch size	10

B. Task 1: 1-object pickup

This task involves a gripper of mass m^g picking up an object of mass m^o while moving in a 2D plane. The independent dynamics of the gripper and object can be written as,

$$\mathbf{x}_{t+1}^g = \mathbf{A}_t \mathbf{x}_t^g + \mathbf{B}_t^g \mathbf{u}_t^g, \quad \mathbf{x}_{t+1}^o = \mathbf{A}_t \mathbf{x}_t^o + \mathbf{B}_t^o \mathbf{u}_t^o, \text{ where}$$

$$\mathbf{x}_t^g = \begin{bmatrix} q_{x_t}^g \\ q_{y_t}^g \\ \dot{q}_{x_t}^g \\ \dot{q}_{y_t}^g \end{bmatrix}, \mathbf{A}_t = \begin{bmatrix} 1 & 0 & dt & 0 \\ 0 & 1 & 0 & dt \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \mathbf{B}_t^g = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ dt/m^g & 0 \\ 0 & dt/m^g \end{bmatrix}$$

$$\mathbf{x}_t^o = \begin{bmatrix} q_{x_t}^o \\ q_{y_t}^o \\ \dot{q}_{x_t}^o \\ \dot{q}_{y_t}^o \end{bmatrix}, \mathbf{B}_t^o = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ dt/m^o & 0 \\ 0 & dt/m^o \end{bmatrix}, \mathbf{u}_t^g = \begin{bmatrix} u_{x_t}^g \\ u_{y_t}^g \end{bmatrix}, \mathbf{u}_t^o = \begin{bmatrix} u_{x_t}^o \\ u_{y_t}^o \end{bmatrix}$$

q_{x_t} , q_{y_t} and \dot{q}_{x_t} , \dot{q}_{y_t} are positions and velocities respectively. It is important to note here that the object is unactuated ($u_{x_t}^o = 0$ at all time t) but when grasped is affected by the force applied by the gripper. We consider two modes: Mode 1 is when gripper is free and mode 2 is when the gripper has grasped the object and is moving with it. The combined dynamics of gripper and object in different modes $\rho_t \in \{1, 2\}$ is as given in (II), where

$$\mathbf{x}_t = \begin{bmatrix} \mathbf{x}_t^g \\ \mathbf{x}_t^o \end{bmatrix}, \mathbf{A}_{1t} = \mathbf{A}_{2t} = \begin{bmatrix} \mathbf{A}_t & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_t \end{bmatrix}$$

$$\mathbf{B}_{1t} = \begin{bmatrix} \mathbf{B}_t^g \\ \mathbf{0} \end{bmatrix}, \mathbf{B}_{2t} = \begin{bmatrix} \mathbf{B}_t^g \\ \mathbf{B}_t^o \end{bmatrix}, \mathbf{u}_t = \begin{bmatrix} \mathbf{u}_t^g \\ \mathbf{u}_t^o \end{bmatrix}$$

$$\mathbf{C}_{\rho_{t-1}, \rho_t} = \rho_{t-1}^\top \rho_t \mathbf{I} + (1 - \rho_{t-1}^\top \rho_t) \mathbf{C}$$

$$C_{1,2} = C_{2,1} = \frac{1}{m^g + m^o} \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & m^g & 0 & 0 & 0 & m^o & 0 \\ 0 & 0 & 0 & m^g & 0 & 0 & 0 & m^o \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & m^g & 0 & 0 & 0 & m^o & 0 \\ 0 & 0 & 0 & m^g & 0 & 0 & 0 & m^o \end{bmatrix}$$

Here, we observe that there are only two parameters that we need to learn to populate the above matrix, $C_1 = \frac{m^g}{m^g + m^o}$ and $C_2 = \frac{m^o}{m^g + m^o}$. Thus, during system identification we learn the four parameters: C_1, C_2, m^g, m^o in addition to the classifier.

Expert trajectories for this task are generated using the following cost parameters for the respective modes in iLQR,

$$Q_1 = Q_2 = \text{diag}([1, 1, 1, 1, 10^{-6}, 10^{-6}, 10^{-6}, 10^{-6}])$$

$$R_1 = \text{diag}([1, 1]), R_2 = \text{diag}([0.5, 0.5])$$

where $\text{diag}()$ represents a diagonal matrix. During training we assume that the cost matrices diagonal.

1) *Task 2: 2-object pickup*: This task considers 3 modes: Mode 1 is when the gripper is free, mode 2 is when the gripper has grasped object 1 and mode 3 is when gripper has grasped object 2. The combined dynamics (II) can be written using:

$$\mathbf{x}_t = \begin{bmatrix} \mathbf{x}_t^g \\ \mathbf{x}_t^{o1} \\ \mathbf{x}_t^{o2} \end{bmatrix}, \mathbf{A}_{1t} = \mathbf{A}_{2t} = \mathbf{A}_{3t} = \begin{bmatrix} \mathbf{A}_t & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_t & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{A}_t \end{bmatrix}$$

$$\mathbf{B}_{1t} = \begin{bmatrix} \mathbf{B}_t^g \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \mathbf{B}_{2t} = \begin{bmatrix} \mathbf{B}_t^g \\ \mathbf{B}_t^{o1} \\ \mathbf{0} \end{bmatrix}, \mathbf{B}_{3t} = \begin{bmatrix} \mathbf{B}_t^g \\ \mathbf{B}_t^{o1} \\ \mathbf{B}_t^{o2} \end{bmatrix}, \mathbf{u}_t = \begin{bmatrix} u_{xt}^g \\ u_{yt}^g \end{bmatrix},$$

In the interest of space, we write the C matrix such that it maps the velocities just before collision to the velocities just after collision, thus ignoring the rows that identically map the position before collision to position after collision.

$$C_{1,2} = C_{2,1} = \frac{1}{m^g + m^{o1}} \begin{bmatrix} m^g & 0 & m^{o1} & 0 & 0 & 0 \\ 0 & m^g & 0 & m^{o1} & 0 & 0 \\ m^g & 0 & m^{o1} & 0 & 0 & 0 \\ 0 & m^g & 0 & m^{o1} & 0 & 0 \\ 0 & 0 & 0 & 0 & m^g + m^{o1} & 0 \\ 0 & 0 & 0 & 0 & 0 & m^g + m^{o1} \end{bmatrix}$$

$$C_{2,3} = \frac{1}{m^g + m^{o1} + m^{o2}} \begin{bmatrix} m^g & 0 & m^{o1} & 0 & m^{o2} & 0 \\ 0 & m^g & 0 & m^{o1} & 0 & m^{o2} \\ m^g & 0 & m^{o1} & 0 & m^{o2} & 0 \\ 0 & m^g & 0 & m^{o1} & 0 & m^{o2} \\ m^g & 0 & m^{o1} & 0 & m^{o2} & 0 \\ 0 & m^g & 0 & m^{o1} & 0 & m^{o2} \end{bmatrix}$$

Thus, to populate each $C_{i,j}$ matrix we need 9 independent parameters. We learn these parameters along with the masses m^g, m^{o1}, m^{o2} and classifier during system identification. It is interesting to note here that we only observe $C_{1,2}$ and $C_{2,3}$ in our demonstrations for this task, hence even though we define all the other $C_{i,j}$ s as training variables they are never learned and do not affect the training time.

Expert trajectories for this task are generated using the following cost parameters for the respective modes in iLQR,

$$Q_1 = Q_2 = Q_3$$

$$= \text{diag}([1, 1, 1, 1, 10^{-6}, 10^{-6}, 10^{-6}, 10^{-6}, 10^{-6}, 10^{-6}, 10^{-6}, 10^{-6}])$$

$$R_1 = \text{diag}([1, 1]), R_2 = R_3 = \text{diag}([0.5, 0.5])$$

where $\text{diag}()$ represents a diagonal matrix. Here again during training we assume that the cost matrices diagonal.

2) *Task 3: Pickup task with ground interactions*: This task captures four modes: Mode 1 is when the gripper moves in free space, mode 2 is when the gripper comes in contact with the table and moves along it, mode 3 is when the gripper comes in contact with the object on the table and pushes it along the table and mode 4 is when the gripper picks up the

object off the table and takes it to the goal location. The dynamics for each mode can be written as,

$$\mathbf{x}_t = \begin{bmatrix} \mathbf{x}_t^g \\ \mathbf{x}_t^o \end{bmatrix}, \mathbf{A}_{1t} = \mathbf{A}_{2t} = \mathbf{A}_{3t} = \mathbf{A}_{4t} = \begin{bmatrix} \mathbf{A}_t & \mathbf{0} \\ \mathbf{0} & \mathbf{A}_t \end{bmatrix}$$

$$\mathbf{B}_{1t} = \mathbf{B}_{2t} = \begin{bmatrix} \mathbf{B}_t^g \\ \mathbf{0} \end{bmatrix}, \mathbf{B}_{3t} = \mathbf{B}_{4t} = \begin{bmatrix} \mathbf{B}_t^g \\ \mathbf{B}_t^o \end{bmatrix}, \mathbf{u}_t = \begin{bmatrix} u_{xt}^g \\ u_{yt}^g \end{bmatrix},$$

Again, writing down the impact dynamics while ignoring the rows that identically map positions before collision to after collision,

$$C_{1,2} = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix},$$

$$C_{2,3} = C_{3,4} = \frac{1}{m^g + m^o} \begin{bmatrix} m^g & 0 & m^o & 0 \\ 0 & m^g & 0 & m^o \\ m^g & 0 & m^o & 0 \\ 0 & m^g & 0 & m^o \end{bmatrix}$$

Thus, each of the above matrices can be populated using 4 independent parameters which we learn along with the masses m^g, m^o and classifier during system identification. The expert trajectories were generated using the following cost parameters for each mode,

$$Q_1 = \text{diag}([1, 10, 1, 1, 10^{-6}, 10^{-6}, 10^{-6}, 10^{-6}])$$

$$Q_2 = Q_3 = \text{diag}([1, 1, 1, 1, 10^{-6}, 10^{-6}, 10^{-6}, 10^{-6}]),$$

$$Q_4 = \text{diag}([1, 1, 0.01, 0.01, 10^{-6}, 10^{-6}, 10^{-6}, 10^{-6}]),$$

$$R_1 = R_2 = \text{diag}([1, 1]), R_3 = \text{diag}([0.5, 0.5]), R_4 = \text{diag}([0.1, 0.1])$$

where $\text{diag}()$ represents a diagonal matrix. Here again during training we assume that the cost matrices diagonal.