

DiffSRL: Learning Dynamical State Representation for Deformable Object Manipulation with Differentiable Simulation

Sirui Chen^{*1}, Yunhao Liu^{*1}, Shang Wen Yao¹, Jialong Li¹, Tingxiang Fan^{1,2}, Jia Pan^{†1,2}

Abstract—Dynamic state representation learning is essential for robot learning. Good latent space that can accurately describe dynamic transition and constraints can significantly accelerate reinforcement learning training as well as reduce motion planning complexity. However, deformable object have very complicated dynamics and is hard to be represented directly by a neural network without any prior physics information. We propose DiffSRL, an end-to-end dynamic state representation learning pipeline that uses differentiable physics engine to teach neural network how to represent high dimensional pointcloud data collected from deformable objects. Our specially designed loss function can guide neural network aware physics constraints and feasibility. We benchmark the performance of our methods as well as other state representation algorithms with multiple downstream tasks on PlasticineLab. Our model demonstrates superior performance most of the time on all tasks. We also demonstrate our model’s performance in real hardware setting with two manipulation tasks on a UR-5 robot arm. The source code is available at <https://github.com/ericcsr/DiffSRL/> and our attached video.

Index Terms—Representation Learning, Deep Learning for Visual Perception, Deep Learning in Grasp and Manipulation

I. INTRODUCTION

Deep neural networks have become a powerful tool in representing high dimensional data using fewer dimensions. Nowadays, its rapid development and application has significantly enhanced the machines to process complex data such as images, sentences and graphs. One application that has recently aroused great research interest is representing real-world objects in dynamic environments from high dimensional point clouds or images for control. Low dimension latent space can benefit learning algorithms by reducing the tunable parameters, accelerating the training, and achieving better control results. The efficiency of this paradigm on rigid body has been demonstrated in recent research [1], [2].

Manuscript received: April, 30, 2022; Accepted July 27, 2022.

This paper was recommended for publication by Editor Jens Kober upon evaluation of the Associate Editor and Reviewers’ comments. This work was supported by HKSAR Research Grants Council (RGC) General Research Fund (GRF) HKU 11202119, 11207818, and the Innovation and Technology Commission of the HKSAR Government under the InnoHK initiative. S. Chen, Y. Liu, S. Yao are also supported by Undergraduate Research Fellowship Program (URFP) of HKU.

^{*} denotes equal contribution. [†] denotes the corresponding author.

¹ S. Chen, Y. Liu, J. Li, S. Yao, T. Fan, J. Pan are with the University of Hong Kong. jpan@cs.hku.hk

² T. Fan and J. Pan are also with the Centre for Garment Production Limited, Hong Kong.

Digital Object Identifier (DOI): see top of this page.




Fig. 1: Similar states but expecting distinct action trajectories.

However, representing commonly seen deformable objects in low dimensions remains a challenge [3]. It may be accounted for by: a) State and dynamic complexity. Deformable objects’ states are usually represented by particle systems such as meshes or point clouds. Their dynamics is usually represented by discretizing complex differential equations on particles. Compared with rigid body, they have much larger state space as well as much more complicate dynamics. b) Correlation complexity. Different particles of the one deformable object can have very different correlations under different deformation [3]. For instance, all points on the surface of a piece of cloth have the same movement when the cloth is translating as a whole. However, when deforming, the cloth is divided into many parts where surface points move independently. c) Topological complexity. Because of different topology under deformation, visually similar states may not have same topology, making it confusing for state representation learning methods to distinguish different states basing on static visual metrics. Fig. 1 shows one example where a rod is to be removed from a piece of plasticine and placed at a target position with minimum energy consumption. The two initial states are similar but have subtle differences: the rod is surrounded by plasticine on the left, while there is a notch on the right. Static representation encoders may overlook their difference because the notch occupies a tiny volume of space. Thus, controller may generate similar actions based on latent state from a static state encoder, but we prefer an algorithm that first drags the rod out (action ①) and then translates it (action ②) in the left case, and directly translates the rod through the notch in the right one. Similar challenges are ubiquitous in soft object manipulation, calling for new state representation learning algorithms, capable of both examining the current state and understanding the dynamics to reflect different futures bifurcated from similar present, which we defined as the “dynamic-awareness.”

Existing state representation learning methods developed

from rigid body systems, such as AutoEncoder [2] and Embed to Control (E2C) [1], lack specific designs for the above challenges related to deformable objects. Newly emerged solutions focusing on deformable objects, such as the Contrastive Forward Model (CFM) [3] or G-Doom [4], as is elaborated in Sec. II, are very data demanding since they fail to utilize physics prior in their training pipelines, leading to a performance bottleneck on more complicate deformable objects other than ropes or cloth under very limited deformation [5]. To improve the performance of the state representation learning on deformable objects, we propose a new pipeline, DiffSRL, that utilizes a differentiable simulator to encode dynamic and constraint-related information.

Our main contributions are two-fold.

- We propose an data efficient end-to-end dynamic state representation learning pipeline utilizing differentiable simulation to capture the multistep dynamics of deformable objects. To the best of our knowledge, this is the first time that state representation learning directly uses the simulation as physics prior as a part of the end-to-end training pipeline. We compared different state representation learning algorithms such as CFM [3] inverse model learning [6], AutoEncoder [2] and E2C [1] with our method on different downstream tasks varying from simple trajectory reconstruction to reinforcement learning algorithm for manipulation. Our method achieves the best performance among others most of the time.
- We design a novel loss function in our proposed method that guides the encoder to be aware of dynamical constraints during training, which can improve the performance of downstream tasks.

We also transfer learned policy from the simulator to the real world. We conducted experiment on a real robot to verify our model against noisy data from a depth camera.

II. RELATED WORK

One of the most explored methods for state representation learning may be the AutoEncoder [2], which trains a neural encoder-decoder network pair to extract features from images to accelerate and enhance the Deep Reinforcement Learning (DRL) by compressing useful information to low-dimension latent state. However it only uses static state information and will fail to distinguish visually similar but dynamically different states [5] as is discussed above. To address this issue, E2C [1] proposes a pipeline including a learned linear dynamical model together with an AutoEncoder. E2C demonstrates its strength in enhancing DRL’s performances in target-reaching problems with image observations, but its linear model cannot describe deformable objects’ complex dynamics well, especially for long horizon. Replacing the E2C’s linear model with a neural network may fail because neural networks with significantly more parameters may be hard to converge with reasonable amount of data and prevent the E2C from achieving desirable outcomes. Alternatively, the inverse-model-based state representation learning [7] is established on the heuristics that good state representation should be able to predict actions leading to the transfer

between two consecutive states. Thus, the encoder can be trained with an action predictor by minimizing the MSE loss between actual and predicted action. The inverse-model-based method has been further applied on model-free RL tasks [8], [9], with more variants and applications elucidated in [5], but existing designs do not consider the complexity of deformable objects either. CFM [3] explores the state representation learning on deformable objects such as cloth and ropes. It utilizes contrastive learning to jointly train a forward model and an encoder to represent images of deformable objects with distinguishable latent states. The trained forward model and the encoder have been used in MPC-based manipulation control. G-Doom [4] applies a similar paradigm to CFM but with a different architecture in terms of using graphs as the latent space representation and graph-dynamic models as the forward model. However, despite their attempts in handling deformable objects, both CFM and G-Doom don’t utilize prior knowledge in physics and require large datasets for relatively simple tasks.

Neural network encoders are the corner stone for a state representation learning method. Designing a neural network architecture for point cloud is challenging due to its unstructured and order-invariant properties. PointNet [10] is a pioneering work of applying deep learning on point cloud data, where the permutation invariance is achieved using point-wise symmetric functions, including point-wise Multi-layer Perceptrons (MLPs) and max-pooling layers. Moreover, point cloud encoders that are derived from the PointNet can tolerate different number of input points, allowing it to be evaluated on point clouds with different densities from that of training data. Thus, we use the PointNet as the backbone of our encoder, and utilizes point clouds extracted from the PlasticineLab MPM simulation as the training dataset.

Recent advances of physics-based simulation, such as Material Point Method (MPM) [11], allow us to accurately and efficiently simulate deformable objects such as cloth and plasticine. Meanwhile, a new branch of physics simulator, the differentiable simulator [12]–[14], make each simulation step differentiable and can be integrated as a layer in deep neural networks, making end-to-end training possible. These simulators provide us with an opportunity to directly introduce physics prior into training pipelines to overcome the difficulty of capturing complex dynamics with a neural net faced by E2C, CFM and G-DOOM. In end-to-end training, gradients from the upstream neural networks are propagated through physics engines and affect network parameters accordingly. For instance, ChainQueen [15] firstly introduces differentiable MPM method to simulate soft robots and build simple controllers. DiffTaichi [14] and PlasticineLab [16] extended the differentiable MPM to simulate liquid and plasticine using the Taichi language [17].

III. PROBLEM FORMULATION

A general dynamic system consists of three components: a state space \mathcal{S} , an action space \mathcal{A} , and a state transfer function:

$$s_{t+1} = f_{\text{sim}}(s_t, a_t), \quad s_t, s_{t+1} \in \mathcal{S}; a_t \in \mathcal{A}. \quad (1)$$

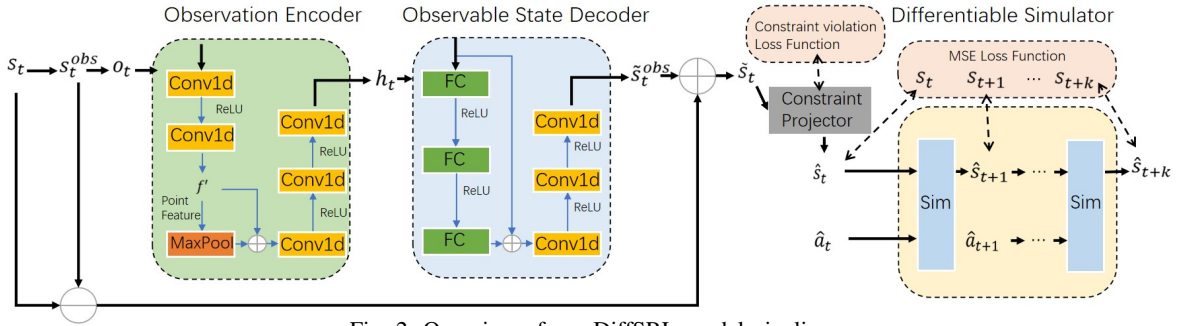


Fig. 2: Overview of our DiffSRL model pipeline.

The state s_t is obtained from the differentiable simulator which contains physical properties such as coordinates and velocities of particles. To mimic real sensor's observation, the observation o_t contains only the positions of a subset of the particles (we call position of all particles as observable state s_t^{obs} , so $o_t \in s_t^{\text{obs}}$). Imperceptible information such as velocities and internal forces are discarded from observation. Our goal is to learn a neural-network encoder f^θ capable of computing a low dimensional latent state h_t from observation: $h_t = f^\theta(o_t)$, where h_t needs to capture sufficient information related to dynamics from the present observation so that it can be used to measure whether two latent states are dynamically similar or not. In particular, two latent states h_t and h'_t are regarded as dynamically equivalent if given an arbitrary action sequence $a_t, a_{t+1}, \dots, a_{t+k}$, their future states s_{t+k}, s'_{t+k} are the same.

IV. APPROACH

The framework of the DiffSRL consists of three main components: 1) a state-observing AutoEncoder which compresses the observation o_t to a latent state h_t and reconstructs the observable state s_t^{obs} from the latent state; 2) a constraint regulator which regulates the states against system dynamical constraints as well as compute additional constraint violation loss to regulate the autoencoder; 3) a differentiable simulator to roll out the trajectory starting from the decoded state. The rolling-out process is differentiable so that gradients can be propagated from the end of the trajectory all the way back to the encoder. The overall pipeline of DiffSRL is shown in Fig. 2.

In this section, we will first introduce the background related to our design, then move on to details of the three major components and our loss function. Finally, we illustrate more details on implementing of our framework on real robots.

A. Preliminaries

Particle-In-Cell paradigm: Lagrangian methods and Euler methods are two major branches of simulation algorithms [11]. The former carries physical properties with particles and is convenient for numerical integration, while the latter maintains physical properties on fixed grids and can be used for fast and interactive simulation. The Particle-In-Cell (PIC) method combines both the methods and is commonly applied in fluid and deformable-object simulation. The deformable-object simulation in the PlasticineLab is based on the Material Point Method [11], a variant of Particle-In-Cell. Our constraint

regulator utilizes the PIC paradigm to detect and resolve penetration efficiently.

Distance Metric on Point Clouds: Point clouds usually are expressed as unordered sets. Hence, differences between two point clouds cannot be measured by common metrics such as MSE or MAE. The Earth Mover's Distance (EMD) [18] is adopted in the DiffSRL to compare two point clouds in terms of the minimum total pairwise distance. The EMD can be expressed as:

$$\text{Dist}_{\text{emd}}(A, B) = \min_{\phi: A \rightarrow B} \sum_{a \in A} d(a, \phi(a)) \quad (2)$$

where A and B are two unordered point sets, $d(\cdot, \cdot)$ is the distance between two sets, and $\phi(\cdot)$ is the optimal correspondence between a and b which is computed iteratively by the auction algorithm. The EMD has been made differentiable thanks to works such as [18] in 3D vision.

B. State Observation Autoencoder

To obtain the latent states of a dynamic system, a deep neural network with learnable parameters θ_{encoder} is used as the encoder, where the specific architecture depends on the type of observations. For instance, the MLP is applied on ordered vector observation, Convolution Neural Network (CNN) on images, and permutation invariant encoders such as PointNet [10] on unordered sets including point clouds. Meanwhile, a decoder with parameters θ_{decoder} is trained simultaneously to recover information from latent states. Formally, the AutoEncoder can be expressed as:

$$h_t = f^{\theta_{\text{encoder}}}(o_t), \quad \tilde{s}_t^{\text{obs}} = f^{\theta_{\text{decoder}}}(h_t) \quad (3)$$

where o_t, h_t are as defined in Sec. III and \tilde{s}_t^{obs} is the reconstructed observation from h_t .

It deserves notice here that since the observation o_t is selected from only the observable part in state s_t , the outcome from the decoder, \tilde{s}_t^{obs} , is expected to contain no imperceptible data as well. Thus, before this outcome enters the rest components, the imperceptible information needs to be appended back to \tilde{s}_t^{obs} to formulate the \tilde{s}_t . This procedure can be expressed as:

$$\tilde{s}_t = s_t \ominus s_t^{\text{obs}} \oplus \tilde{s}_t^{\text{obs}} \quad (4)$$

where s_t^{obs} denotes the observable part of a state. Hence, $s_t \ominus s_t^{\text{obs}}$ gives the imperceptible part of the state, and the \oplus operation adds the reconstructed \tilde{s}_t^{obs} back.

C. Differentiable Constraint Projector

The reconstructed state \tilde{s}_t may violate hard dynamics constraints, such as non-interpenetration constraints, joint limit constraints, and continuity constraints within continuous deformable objects. Constraint violation usually causes meaningless gradient, significant numerical errors or even simulator failure. Hence, it is necessary to make sure the input to the differentiable simulator is always valid as well as provide incentive for the AutoEncoder's to generate feasible state throughout the training procedure. Therefore, a constraints regulator is designed to project decoder output state to the closest the feasible state and provide incentive for the AutoEncoder to stay away from unrealistic states. Formally, the reconstructed state, the constraint regulator loss, and the projected state are computed as:

$$\mathcal{L}_{\text{constraint}} = \min_{s \in \mathcal{S}} \text{Dist}_{\text{emd}}(s, \tilde{s}_t), \quad (5)$$

$$\hat{s}_t = \arg \min_{s \in \mathcal{S}} \text{Dist}_{\text{emd}}(s, \tilde{s}_t), \quad (6)$$

where the distance Dist is the aforementioned EMD Eq. (2) between point clouds. We will explain details of how to compute the closest feasible state in \mathcal{V} as well as 1

D. Differentiable Simulator and loss design

By using the reconstructed state from decoder, the simulation in Eq. (1) then becomes

$$\hat{s}_{t+1} = f_{\text{sim}}(\hat{s}_t, a_t) \quad (7)$$

where \hat{s}_t is output state of decoder after constraint projection, from which the simulator f_{sim} executes the input action a_t and reaches the corresponding result state \hat{s}_{t+1} . When applied to a trajectory of length k , a sequence $\hat{s}_{t:t+k}$ can be thereafter retrieved starting from \hat{s}_t , along an action sequence $a_t, a_{t+1}, \dots, a_{t+k}$.

The model is trained with both the constraints violation loss $\mathcal{L}_{\text{constraint}}$ (Eq. (5)) and the multistep reconstruction loss $\mathcal{L}_{\text{multi-step}}$ defined as

$$\mathcal{L}_{\text{multi-step}} = \sum_{i=1}^k \gamma^i \text{Dist}_{\text{emd}}(s_{t+i}, \hat{s}_{t+i}), \quad (8)$$

The multistep reconstruction loss penalizes the distance between each state s_{t+i} in the reference trajectory from the dataset and its correspondence in the rolling-out trajectory $\hat{s}_{t:t+k}$. A decaying factor γ is used to mitigate the gradient instability effect along backpropagation through a long horizon. In this way, the total loss to be optimized is

$$\mathcal{L}_{\text{total}} = (1 - \beta)\mathcal{L}_{\text{multi-step}} + \beta\mathcal{L}_{\text{constraints}} \quad (9)$$

where the weight factor β is introduced to trade-off among different loss terms during training. Initially, the decoded state violates constraints significantly, so β is set to 1 to encourage the AutoEncoder to respect the physical constrains. Progressively, β decays exponentially at rate λ so that the multistep reconstruction loss dominates.

E. Implementation Details

The MPM of the PlasticineLab provides a direct access to the particles that composes the deformable objects. Therefore, as is declared in Sec. III, we set each observation o_t as a subset of the positions of all particles, which can be naturally considered as an unordered point cloud. Hence, a variant of Point Completion Network (PCN) [19] which is an extension of the PointNet ideology, is used to implement the AutoEncoder.

As for the constraint projector, the MPM simulation disallows penetration between particles and rigid bodies and also requires smooth velocity fields within continuous materials. Hence, our projector is constructed in two stages: 1) We use signed distance field (SDF) with MPM to efficiently detect penetration. For those particles p_{collide} who have negative signed distance and reside inside rigid bodies, we resolve the penetration by finding the minimum distance $d_{p_{\text{collide}}}$ to exit the colliding object. For simplicity, it is assumed that each rigid-body object is a geometric primitive and has an analytical collision detection routine so that the minimum exit distance can be solved efficiently using a simple linear programming solver. The loss function for this objective is defined as the sum of all $d_{p_{\text{collide}}}$. 2) We use auction algorithm [20], a GPU parallelizable point correspondence matching algorithm, to find the best pairwise matches between particles in reference state and those in the reconstructed state after collision resolution in 1). The step is necessary because, as is emphasized in Sec. IV-B, the imperceptible information such as velocity and shearing force of each particle is still missing in the decoder's outputs \tilde{s}_t^{obs} . The information needs to be consistent with field properties maintained by the grid to execute the MPM properly. After the auction algorithm establishes pairwise relationships, each ground-true particle's properties are exerted to its match, which fulfills the \oplus notation in 4. The details are described in Algorithm 1 and also illustrated in Fig. 3. The pairwise distance between the reference state and projected state is also added as part of constraint satisfaction loss.

F. Closing the Sim2Real Gap

In the simulator, the observable part of states s_t^{obs} are in the form of dense point clouds. However, a real-world sensor, such as a depth camera and a LiDAR, can only provide a thin-shell of noisy point cloud of an object's surface. Uniformly

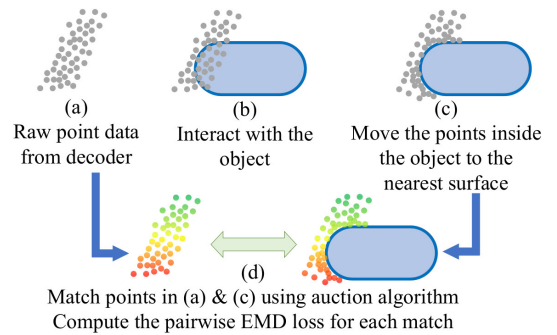


Fig. 3: The collision-free and smooth-velocity-field constraint projector, where the dots and the capsule represent the point clouds and the rigid-body object which may collide with them, respectively.

Algorithm 1 Constraints Projector

Input: Rigid bodies \mathcal{R} , the ground-truth particle observable state s_t^{obs} (i.e., position of particles), and decoded particles position \hat{s}_t^{obs}

Initialize: Grids to store interpenetrating rigid bodies information $c_{\text{grid}}^{\text{col}}$, loss of penetration $l_{\text{penetration}}$, and loss reconstruction l_{rec} .

$\triangleleft \hat{s}_t^{\text{obs, no-penetration}} \leftarrow \text{ResolvePenetration}(\mathcal{R}, \hat{s}_t^{\text{obs}})$

\triangleleft Compute grid mass using particle-in-cell

$m_{\text{grid}} \leftarrow \text{ComputeGridMass}(\hat{s}_t^{\text{obs}})$

for r in \mathcal{R} **do**

\triangleleft Check penetration using signed distance field

$c_{\text{grid}}^{\text{col}} \leftarrow \text{SDF}(r, m_{\text{grid}})$

end for

\triangleleft Compute Interpenetration related information of each particle to all rigid bodies

for p in $\hat{s}_t^{\text{obs no-penetration}}$ **do**

if p.grid is not empty **then**

\triangleleft Solve linear programming problem to find minimum displacement for a particle to stay out of penetrating rigid body

 p.pos $\leftarrow \text{LPSolver}(\mathcal{R}, p)$

\triangleleft Sum up minimum displacement costs as loss

$l_{\text{penetration}} += \text{LPCost}(r, p)$

end if

end for

$\triangleleft \hat{s}_t^{\text{obs}} \leftarrow \text{ResolveMissMatch}(\hat{s}_t^{\text{obs no-penetration}})$

order $\leftarrow \text{AuctionAlgorithm}(\hat{s}_t^{\text{obs}}, \hat{s}_t^{\text{obs}})$

$\hat{s}_t^{\text{obs}} \leftarrow \text{perm}(\hat{s}_t^{\text{obs no-penetration}}, \text{perm}(\text{order}))$

return \hat{s}_t^{obs}

sampling new points within the interior of thin-shell can produce a dense point clouds filling the object’s volume, but such point clouds may observe a different distribution compared with those in the simulator. Thus, besides uniformly sampling new points within the thin-shell when the algorithm is tested on real robot, we augment the dataset by adding a random normal noise to point positions and randomly removing 10% - 50% particles from the dense point cloud during the training in the simulator.

V. EXPERIMENTS AND RESULTS

To assess the effectiveness of the DiffSRL, the model is evaluated with four experiments: 1) trajectory reconstruction, 2) reward prediction, 3) KNN-EMD, and 3) model-free reinforcement learning on 6 single object manipulation tasks from the PlasticineLab benchmark [16]. The proposed DiffSRL has achieved state-of-the-art performance in these experiments most of the time.

A. Baselines for Comparison

We compared our approach against state-of-the-art method CFM [3] and four commonly used dynamic state representation learning methods: the *E2C* method [1], the *Forward* method [21] that learns a dynamic forward model on latent states, the *Inverse* method [6] that conducts inverse model

	predictive dynamic model	predict action	observation reconstruction	multistep information	physical constraints
DiffSRL (Our)	✓		✓	✓	✓
E2C [1]	✓		✓		
Forward [21]	✓				
Inverse [6]		✓			
CFM [3]	✓				
Autoencoder [2]			✓		

TABLE I: Comparison of different components used in different state representation learning pipelines.

learning for action prediction, as well as the *AutoEncoder* method directly reconstructing the state [2]. In Tab. I, we summarize characteristics of different approaches.

Regarding implementation details of these baselines, similar to the DiffSRL, we use EMD [18] as a loss function to train the AutoEncoder in the baseline. The forward model used in the E2C and the Forward $f_{\text{fwd}}^\theta(h_t, a_t)$ is a 3-layer MLP on latent space. The action predictor $\hat{a}_t = f_{\text{inverse}}^\phi(h_t, h_{t+1})$ in the Inverse uses a pair of successive latent states as inputs and tries to infer the action with a 3-layer MLP. The contrastive loss in [3] uses the same formation and negative sample selection method as in the original paper. All MLPs mentioned here use 256 as the size of a hidden layer.

B. Environment selection and Data Collection

We select 6 environments from PlasticineLab that involves single deformable objects. As for *TripleMove*, *table*, *Assemble*, these experiment involves multiple deformable objects and require special design of loss functions as well as adaptation. We left handling multi-objects as future work. For each environment, we collect 6,000 trajectories of 8-step length from simulator. 30% of data are sampled from random policy, while others are collected when optimizing trajectories for manipulation. The reason for doing so is that random actions frequently cause manipulators detaching from the deformable objects too early to induce sufficient deformation. We use the ratio of 8:1:1 to split collected samples into the training, validation and testing sets.

C. Comparison in Trajectory Reconstruction and Reward Prediction

The trajectory reconstruction error and the reward prediction error are two straightforward metrics to reflect the encoder’s capability of dynamics awareness. The trajectory reconstruction error measure the similarity between the reference trajectory and the trajectory roll-out from the reconstructed state given the same action sequence, which is computed by accumulating the earth mover’s distance [18], i.e., $\sum_{s_{t:t+k}} D_{\text{EMD}}(s_t, \hat{s}_t)$, between each pair of states in both trajectories. Some baseline methods such as *CFM*, *Forward* and *Inverse* only train the encoder without decoder and thus cannot reconstruct states. Thus, we trained a decoder for each fixed encoder using the AutoEncoder training pipeline until convergence. Tab. II shows that DiffSRL achieves the best performance among all models.

The reward prediction error evaluates whether a encoded latent state can capture enough information from a state to

Datasets	Rope			Chopsticks			Torus			Writer			Rollingpin			Pinch		
	EMD	MSE	KNN	EMD	MSE	KNN	EMD	MSE	KNN	EMD	MSE	KNN	EMD	MSE	KNN	EMD	MSE	KNN
AutoEncoder	0.048	0.236	0.830	0.044	1.102	<u>0.350</u>	0.049	<u>0.011</u>	<u>0.021</u>	0.057	0.005	0.169	0.037	0.019	0.103	0.031	0.005	<u>0.071</u>
Inverse	0.042	0.276	1.043	0.048	1.190	<u>0.449</u>	0.041	<u>0.099</u>	<u>0.028</u>	0.067	0.008	0.017	0.032	0.093	0.158	0.026	0.046	<u>0.131</u>
Forward	0.083	5.435	<u>0.816</u>	0.124	6.599	0.357	0.192	0.085	0.021	0.214	0.007	<u>0.016</u>	0.091	0.738	0.112	0.512	0.039	0.075
E2C	0.048	<u>0.223</u>	0.915	0.064	1.244	0.359	0.039	0.098	0.026	0.058	0.008	0.015	0.034	0.032	0.153	0.027	0.043	0.166
CFM	<u>0.037</u>	<u>0.236</u>	1.137	0.037	<u>1.010</u>	0.491	0.029	0.012	0.032	0.078	<u>0.004</u>	0.036	0.038	<u>0.015</u>	0.124	0.026	<u>0.004</u>	0.161
DiffSRL	0.027	0.218	0.797	0.031	0.840	0.324	<u>0.036</u>	0.009	0.020	0.044	0.003	<u>0.016</u>	0.021	0.013	<u>0.103</u>	0.016	0.003	0.071
DiffSRL ⁻	0.039	0.292	0.821	<u>0.036</u>	1.226	0.376	0.038	0.012	0.023	<u>0.051</u>	0.006	0.018	<u>0.028</u>	0.017	0.121	<u>0.021</u>	0.004	0.091

TABLE II: Comparison of the average of trajectory rollout earth mover’s distance and mean squared error of reward prediction as well as KNN-MEMD. The smaller the MSE and EMD, the better the model. The best result is highlighted in **bold** and the second best is highlighted with underline. DiffSRL⁻ is our proposed method trained without non-penetration constraint projector.

accurately predict reward for a given task. For each task, we trained a 3-layer MLP for latent states obtained from each encoder to predict rewards from latent states. The output of MLP is reward related to the pose of deformable object. The error is measured using mean square error: $\frac{1}{|\mathcal{D}|} \sum_{s \in |\mathcal{D}|} \|\hat{r}_{\text{pred}} - r_{\text{real}}\|^2$. As summarized in Tab. II, our model achieves the best performance on most of datasets. One thing worth noting is that the *Forward* model observes the worst performance on both metrics among all models, probably because it tends to be trapped in a trivial minimum as mentioned in [5].




Fig. 4: The architecture of policy network that we used in model-based policy optimization and model-free reinforcement learning.

D. KNN-MEMD

KNN-MSE [22] has been used to as a metric for static state representation learning. It can evaluate whether neighbors on latent space are actually close to each other on state space. For dynamic state representation, we need not only check whether two states from a neighborhood are similar in current time stamp, but also future states driven by the same action sequence. Moreover, since two states are two point clouds in our case, we replace distance metric from Mean Square Error(MSE) to Mean Earth Mover’s Distance(MEMD). We can expect that dynamical equivalent states will have smaller KNN-MEMD. Our novel metrics is evaluated as follow:

$$\text{KNN-MEMD}(h) = \frac{1}{K} \sum_{h' \in \text{knn}(h, K)} \frac{1}{H} \sum_t^H \text{EMD}(s_t, s'_t) \quad (10)$$

H is the rollout horizon, K is number of neighbors, h, h' are latent states and s, s' are states. Π

E. Model-Free Reinforcement Learning (MFRL)

To study the effectiveness of state representation learning on various complex downstream tasks, we evaluated encoders trained from different methods on training model free reinforcement learning. We uses TD3, a state-of-the-art off policy method to train policies based on latent state observation. The

trained policy is expected to control manipulators to move the plasticine to different target positions within finite steps. Detailed target and reward design are the same as described in [16]. We use latent states for both state value function and policy function. The architecture of the policy network is shown in Fig. 4, where the encoder’s weights are fixed during policy learning. For comparison, we added another baseline: the naive down-sampling based method provided by PlasticineLab, which samples 400 points from each point cloud¹. The policy network is initialized as a simple MLP, and its inputs are obtained by concatenating features of these 400 particles. We evaluated performance using the per-epoch average rewards. Each experiment is repeated for five times, and the means and standard deviations are plotted in Fig. 5. From the first row we can see that most state representation learning methods can help policy improving faster and achieve better final rewards compared with the down-sampling based method. Moreover, our approach achieves the best result among all state representation learning methods, including the *AutoEncoder* that is trained using a similar loss function. Some keyframes from the trajectory are shown in Fig. 6.

F. Ablation Study

To evaluate the effectiveness of the constraint projector, we conduct an ablation study by removing non-penetration projector during the training. We compare its result with our method both quantitatively on different metrics as shown in II and qualitatively using MFRL training as in 8. Regarding the trajectory reconstruction errors, the reward prediction error as well as KNN-MEMD, it shows that removing the non-penetration projector will result in performance degradation. We also tried to remove the smoothness projector during training, but this will lead to numerical issues in the differentiable simulator because it violate MPM’s requirement that particles in the same grid have consistent properties. Furthermore, the encoders without non-penetration constraints projector shows worse result in reinforcement learning. Therefore, it can show that constraints projector is essential for improving performance.

G. Robustness Analysis

The number of points observed by a sensor may vary significantly in real applications. Since the PCN-based point

¹The original down-sampling based method from PlasticineLab is not feasible in practice since it requires accurately tracking 400 points within the deformable object, which is impractical.

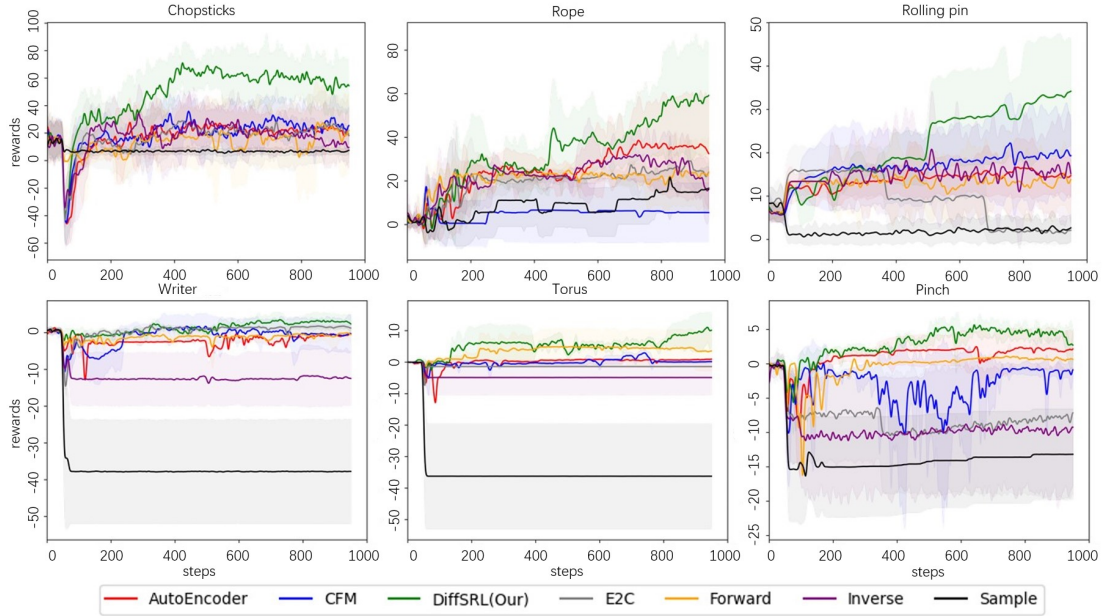


Fig. 5: TD3 Reward in different environments.

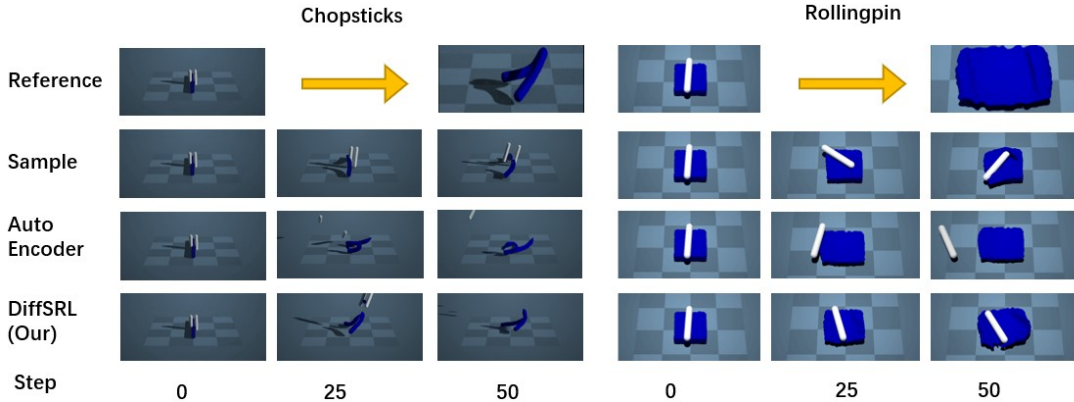


Fig. 6: Keyframes of different approaches in Model-Free Reinforcement Learning (MFRL) task.




Fig. 7: MFRL reward in Chopsticks when using different number of point as encoder input. Averaging over 5 experiments for each input particle number.




Fig. 8: Ablation Study: without non-penetration projector

cloud encoder can be directly applied with different point number without re-training, we used various numbers of observed points to test the robustness of our model. We trained a DiffSRL model using point clouds consisting of 8,000 points and then investigated the deployment performance using 6,000, 4,000, 2,000 particles by computing the reward curves when training model free reinforcement learning from latent states. As shown in Fig. 7, although using fewer observable particles (from 8,000 to 6,000) can decrease the performance, further decreasing of the particle number (from 6,000 to 1,000) will not significantly downgrade the performance, which demonstrates that our method is relatively robust against disturbance in observation. Only we set number of observable particles to 500, can we observe significant decay in downstream task's performance.

H. Hyperparameters Setting

Tab. III summarizes the hyperparameters used in our model. One important hyperparameter of our method is k , the number of trajectory rollout steps. We choose k by evaluating the 10-step trajectory reconstruction loss on the validation set, with the result shown in Fig. 9. For both environments,




Fig. 9: 10-step EMD loss on both rope and chopstick when using different of training steps.

Parameters	Meaning	Value
N	Number of particles	8192
I	ICP iterations	3000
d_{latent}	Latent Space dimension	1024
α	Learning rate	1e-5
γ	weight decay rate in 8	0.99
β	Initial weight between in 9	0.99
λ	Decay rate of β per epoch	0.9
E	Number of epochs	20

TABLE III: Hyperparameters




Fig. 10: Simulated and real-robot experiments for Rope and Chopsticks.

$k = 7$ achieves the best performance. Notice that the loss curve demonstrates non-monotonic behavior as we increase the training step size, and the possible reason is that as the step size increases, the gradient back-propagation through time tends to be noisier due to numerical issues, and the optimization landscape becomes more wiggly. We leave a further investigation as future work.

I. Hardware Experiment

We deployed both policies trained with MFRL on a UR5 robot with a Kinect-v2 depth camera. Both tasks require the a robotic arm to manipulate a plasticine on a plane as shown in Fig. 10. We first segment the plasticine from the background using color. Next, to obtain a dense point cloud for the volume taken by the plasticine, we use the points on the top of the plasticine’s surface as seeds. For each seed we sampled more points below it to infill the volume. Fig. 10 also shows the start and terminal states of these two tasks both in simulation and real world. As demonstrated, our method can work well on real robots. We found that the occlusion could be an issue and thus we designed the end-effector to be thin and tall to minimize the possible occlusion. How to handle the partial observation will be left as future work.

VI. CONCLUSION AND FUTURE WORK

We presented a novel dynamic state representation learning model for deformable objects. Moreover, we evaluated our

model with multiple tasks on soft-body manipulation in PlasticineLab [16] and compared its performance with other state-of-the-art models. Meanwhile, the applicability of the method on real-world scenarios has been examined using the robotic arms. One possible follow-up might include the differentiable rendering method [23] in our pipeline, which enables using more accessible images as observation while maintaining the overall differentiability.

REFERENCES

- [1] M. Watter, J. T. Springenberg, J. Boedecker, and M. A. Riedmiller, “Embed to control: A locally linear latent dynamics model for control from raw images,” in *NeurIPS*, 2015.
- [2] C. Finn, X. Y. Tan, Y. Duan, T. Darrell, S. Levine, and P. Abbeel, “Deep spatial autoencoders for visuomotor learning,” in *ICRA*, 2016.
- [3] W. Yan, A. Vangipuram, P. Abbeel, and L. Pinto, “Learning predictive representations for deformable objects using contrastive estimation,” in *CoRL*, 2020.
- [4] X. Ma, D. Hsu, and W. S. Lee, “Learning latent graph dynamics for deformable object manipulation,” in *CoRL*, 2021.
- [5] T. Lesort, N. D. Rodríguez, J. Goudou, and D. Filliat, “State representation learning for control: An overview,” *Neural Networks*, 2018.
- [6] “Learning to poke by poking: Experiential learning of intuitive physics,” in *NeurIPS*, 2016.
- [7] W. Duan and K. Jens, “Learning state representations for robotic control: Information disentangling and multi-modal learning,” Master’s thesis, TU Delft, 2017.
- [8] E. Shelhamer, P. Mahmoudieh, M. Argus, and T. Darrell, “Loss is its own reward: Self-supervision for reinforcement learning,” in *ICLR*, 2017.
- [9] D. Pathak, P. Agrawal, A. A. Efros, and T. Darrell, “Curiosity-driven exploration by self-supervised prediction,” in *CVPR Workshops*, 2017.
- [10] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, “Pointnet: Deep learning on point sets for 3d classification and segmentation,” in *CVPR*, 2017.
- [11] C. Jiang, C. A. Schroeder, J. Teran, A. Stomakhin, and A. Selle, “The material point method for simulating continuum materials,” in *SIGGRAPH*, 2016.
- [12] K. Werling, D. Omens, J. Lee, I. Exarchos, and C. K. Liu, “Fast and feature-complete differentiable physics engine for articulated rigid bodies with contact constraints,” in *RSS*, 2021.
- [13] J. Xu, T. Chen, L. Zlokapa, M. Foshey, W. Matusik, S. Sueda, and P. Agrawal, “An end-to-end differentiable framework for contact-aware robot design,” in *RSS*, 2021.
- [14] Y. Hu, L. Anderson, T. Li, Q. Sun, N. Carr, J. Ragan-Kelley, and F. Durand, “DiffTaichi: Differentiable programming for physical simulation,” in *ICLR*, 2020.
- [15] Y. Hu, J. Liu, A. Spielberg, J. B. Tenenbaum, W. T. Freeman, J. Wu, D. Rus, and W. Matusik, “Chainqueen: A real-time differentiable physical simulator for soft robotics,” in *ICRA*, 2019.
- [16] Z. Huang, Y. Hu, T. Du, S. Zhou, H. Su, J. B. Tenenbaum, and C. Gan, “Plasticinelab: A soft-body manipulation benchmark with differentiable physics,” in *ICLR*, 2021.
- [17] Y. Hu, T. Li, L. Anderson, J. Ragan-Kelley, and F. Durand, “Taichi: a language for high-performance computation on spatially sparse data structures,” *ACM Trans. Graph.*, 2019.
- [18] Y. Rubner, C. Tomasi, and L. J. Guibas, “A metric for distributions with applications to image databases,” in *ICCV*, 1998.
- [19] W. Yuan, T. Khot, D. Held, C. Mertz, and M. Hebert, “PCN: point completion network,” in *3DV*, 2018.
- [20] D. P. Bertsekas, “A distributed algorithm for the assignment problem,” MIT, Tech. Rep., 1979.
- [21] L. Kaiser, M. Babaeizadeh, P. Milos, B. Osinski, R. H. Campbell, K. Czechowski, D. Erhan, C. Finn, P. Kozakowski, S. Levine, A. Mohiuddin, R. Sepassi, G. Tucker, and H. Michalewski, “Model based reinforcement learning for atari,” in *ICLR*, 2020.
- [22] P. Sermanet, C. Lynch, J. Hsu, and S. Levine, “Time-contrastive networks: Self-supervised learning from multi-view observation,” in *CVPR Workshops*, 2017.
- [23] H. Kato, D. Beker, M. Morariu, T. Ando, T. Matsuoka, W. Kehl, and A. Gaidon, “Differentiable rendering: A survey,” *arxiv 2006.12057*, 2020.