

# Learning Predictive Representations for Deformable Objects Using Contrastive Estimation

Wilson Yan<sup>1,†</sup>, Ashwin Vangipuram<sup>1</sup>, Pieter Abbeel<sup>1</sup>, and Lerrel Pinto<sup>1,2</sup>



Fig. 1: Example trajectories using our contrastive forward model for rope and cloth manipulation. The top two rows show rope manipulation from different start states to different goal states, while the bottom two rows show cloth manipulation using different colored cloths. Note that in the last row, the robot is manipulating a white cloth, but our method is able to still use a blue cloth as the goal image.

**Abstract**—Using visual model-based learning for deformable object manipulation is challenging due to difficulties in learning plannable visual representations along with complex dynamic models. In this work, we propose a new learning framework that jointly optimizes both the visual representation model and the dynamics model using contrastive estimation. Using simulation data collected by randomly perturbing deformable objects on a table, we learn latent dynamics models for these objects in an offline fashion. Then, using the learned models, we use simple model-based planning to solve challenging deformable object manipulation tasks such as spreading ropes and cloths. Experimentally, we show substantial improvements in performance over standard model-based learning techniques across our rope and cloth manipulation suite. Finally, we transfer our visual manipulation policies trained on data purely collected in simulation to a real PR2 robot through domain randomization.

## I. INTRODUCTION

Robotic manipulation of rigid objects has received significant interest over the last few decades, from grasping novel objects in clutter [28, 25, 47, 39, 12] to dexterous in-hand manipulation [22, 2, 59]. However, the objects we interact within our daily lives are not always rigid. From putting on clothes to packing a shopping bag, we constantly need to manipulate objects that deform. Even seemingly rigid

objects like metal wires significantly deform during everyday interactions. As a result, there has been a growing interest in algorithms that can tackle deformable object manipulation [54, 15, 42, 43, 45, 58, 46, 29, 50].

Deformable object manipulation presents two key challenges for robots. First, unlike rigid objects, there is no direct representation of the state. Consider the manipulation problem in Figure 1, where the robot needs to straighten a rope from a start configuration to any goal configuration. How does one track the shape of the rope? This lack of a canonical state often limits representations to discrete approximations [3]. Second, the dynamics of deformable objects are complex and non-linear [9]. Due to microscopic interactions within the object, even simple objects can exhibit complex and unpredictable behavior [38], which makes modeling and performing traditional task and motion planning with such deformable objects difficult.

One class of techniques that circumvents the challenges in state estimation and dynamics modeling is image-based model-free learning [13, 44, 26]. For instance, Matas et al. [30], Seita et al. [46], Wu et al. [58] use model-free methods in simulation for several difficult cloth manipulation tasks. However, without expert demonstrations, model-free learning is notoriously inefficient [7], and often needs millions of samples to learn from. This challenge is further exacerbated in the multi-task learning framework, where the robot needs

<sup>1</sup>Department of EECS, University of California, Berkeley.

<sup>2</sup>Department of Computer Science, New York University.

<sup>†</sup>Correspondence to wilson1.yan@berkeley.edu

to learn to reach multiple goals.

Model-based techniques, on the other hand, have shown promise in sample-efficient learning [57, 4, 35]. However, using such model-based learning techniques for deformable objects necessitates tackling the challenges of state representation and dynamics modeling head-on. So how does one learn models given high-dimensional observations and complex underlying dynamics? Some approaches take a direct approach to learning complex dynamics models through pixel-space [19, 8]. Another approach, by Agrawal et al. [1], Nair et al. [36], learns forward dynamics models in conjunction with inverse dynamic models for manipulating deformable objects. However, during robotic execution, only the inverse model is used. Other model-based approaches such as Wang et al. [56] train Causal InfoGANs [23, 6] to both extract visual representations and forward models, and use the learned forward models for planning. However, these techniques are not robust due to training instabilities associated with GANs [49].

In this paper, we introduce a new visual model-based framework that uses contrastive optimization to jointly learn both the underlying visual latent representations and the dynamics models for deformable objects. We hypothesize that using contrastive methods for model-based learning achieves better generalization and latent space structure do to its inherent information maximization objective. We re-frame the objective introduced in contrastive predictive coding [37] to allow for learning effective model dynamics and latent representations. Once the latent models for representations and dynamics are learned across offline random interactions, we use standard model predictive control (MPC) with one-step predictions to manipulate deformable objects to desired visual goal configurations. Given this controller, we empirically demonstrate substantial improvements over standard model-based learning approaches across multi-goal rope and cloth spreading manipulation tasks. Videos of our real robot runs and reference code can be found on the project website: <https://sites.google.com/view/contrastive-predictive-model>.

In summary, we present three key contributions in this paper: (a) We propose a contrastive predictive modeling approach to model learning that is compatible with model predictive control. To our knowledge, this is the first use of contrastive estimation for model-based learning. (b) We demonstrate substantial improvements in multi-task deformable object manipulation over other model learning approaches. (c) We show the applicability of our method to real robot rope and cloth manipulation tasks by using sim-to-real transfer without additional real-world training data.

## II. RELATED WORK

### A. Deformable Object Manipulation

There has been a substantial amount of prior work in the area of robotic manipulation of deformable objects. A detailed survey of past work can be found in Khalil and Payeur [20], Henrich and Wörn [15].

A standard approach to tackling deformable object manipulation is to use deformable object simulations with planning

methods [18]. Past work in this domain has focused on simple linear deformable objects [41, 55, 34], creating better simulations [40], and faster planning [11]. However, the large number of states for deformable objects makes it difficult to plan correctly while being computationally efficient.

Instead of directly planning on the full dynamics, some prior research has focused on planning on simpler approximations, by using local controllers to handle the actual complex dynamics. One approach to using local controllers is model-based servoing [48, 54], where the end-effector is controlled to a goal location instead of explicit planning. However, since the controller is optimized over simple dynamics, it often gets stuck in local minima with more complex dynamics [32]. To solve this, several works [3, 31] have proposed Jacobian controllers that do not need explicit models, while [17, 16] have proposed learning-based techniques for servoing. We note that our proposed work on learning latent dynamics models is compatible with several of these model-based optimization techniques.

### B. Contrastive Prediction

Learning good representations remains a difficult challenge in deformable object manipulation. There has been a large amount of prior work on contrastive predictive methods to learn better representations of data. Word2Vec [33] optimizes a contrastive loss to demonstrate semantic and syntactic structure in the learned latent space for words. Oord et al. [37] shows that it is possible to learn high-level representations of images, video, and speech data by employing a large number of negative samples. Tian et al. [52] learns high-level representations by encouraging different views of scenes to be embedded close to one another, and further from others through a similarly framed contrastive loss. Recently, SimCLR [5], another contrastive learning framework, achieved state-of-the-art results in self-supervised learning representations, bridging the gap with supervised learning.

## III. CONTRASTIVE FORWARD MODELING (CFM)

In this section, we describe our proposed framework for learning deformable object manipulation: Contrastive Forward Modeling (CFM). We begin by discussing formalism for predictive modeling and contrastive learning. Following that, we discuss our method for learning contrastive predictive models. See Figure 2 for an overview of our training scheme.

### A. Dynamic Predictive Models

For our problem setting, we consider a fully observable environment with observations  $o \in \mathcal{O}$ , actions  $a \in \mathcal{A}$ , and deterministic transition dynamics  $f(o_t, a_t) = o_{t+1}$ . We would like to learn a predictive model  $\hat{f}(o_t, a_t) \approx o_{t+1}$  to approximate the observation of the next timestep. This can be done by directly learning a visual model through pixel space with regression over observation-action-observation tuples [10, 19]. Once we have successfully learned a predictive model, it is natural to use it for planning to reach different desired goal states, for example, different configurations of a rope or

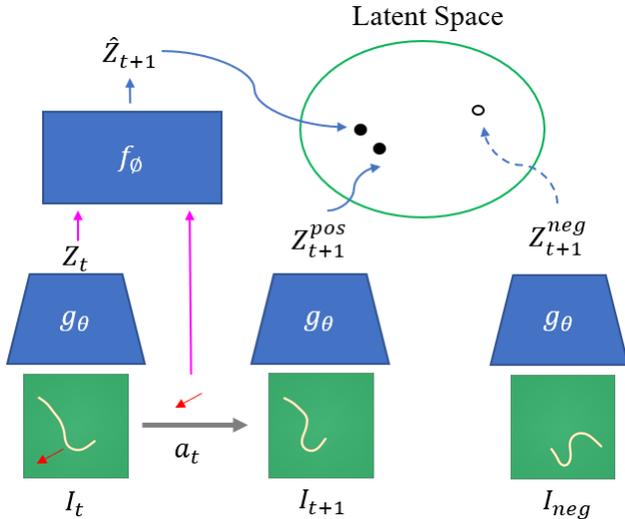


Fig. 2: Overview of our contrastive forward model. Training data consists of (image, next image, action) tuples and we learn the encoder and forward model jointly. The contrastive loss objective brings the positive embedding pairs closer together and the negative embeddings further away.

cloth. However, planning directly through pixel space can be difficult, as pixel-value comparisons between images usually do not necessarily correlate well with their true distances. For example, consider an environment with a ball, where the task is to learn a policy that pushes the ball to the center. If the ball is far from the center, then all predicted next actions using a visual forward model would be equidistant from the goal ball-in-center image when comparing pixels values since there would be no image overlap. Therefore, we consider the framework of planning with in a learned latent space by encoding observations. We learn an encoder  $g_\theta(o_t) = z_t$  to embed our observations into a latent space, coupled with a predictive model in latent space between  $z$ 's, where our learned predictive model is now formulated as  $\hat{f}(z_t, a_t) \approx z_{t+1}$ . In this work, we propose to learn the latent space using a contrastive learning method.

### B. Contrastive Models

In our contrastive learning framework, we jointly learn an encoder  $g_\theta(o_t) = z_t$  and a forward model  $f_\phi(z_t, a_t) \approx z_{t+1}$ . We use the InfoNCE contrastive loss described by Oord et al. [37].

$$\mathcal{L} = -\mathbb{E}_{\mathcal{D}} \left[ \log \frac{h(\hat{z}_{t+1}, z_{t+1})}{\sum_{i=1}^k h(\hat{z}_{t+1}, \tilde{z}_i)} \right] \quad (1)$$

where  $h$  is some similarity function between the computed embeddings from the encoder. The  $\tilde{z}_i$  represents negative samples, which are incorrect embeddings of the next state, and we use  $k$  such negative samples in our loss. The motivation behind this learning objective lies with maximizing mutual information between the predicted encodings and their respective positive samples. Within the embedding space, this

results in the positive sample pairs being aligned together but the negative samples pushed further apart, as seen in Figure 2. Since we are jointly learning a forward model that seeks to minimize  $\|f_\phi(z_t, a_t) - z_{t+1}\|^2$ , we use the similarity function:

$$h(z_1, z_2) = \exp(-\|z_1 - z_2\|^2) \quad (2)$$

where the norm is a  $\ell_2$ -norm. After learning the encoder and dynamics model, we plan using a simple version of Model Predictive Control (MPC), where we sample several actions, run them through the forward model from the current  $z_t$ , and choose the action  $a_t$  that produces  $\hat{z}_{t+1}$  closest (in  $\ell_2$ -distance) to the goal embedding.

## IV. EXPERIMENTAL EVALUATIONS

In this section, we experimentally evaluate our method in various rope and cloth manipulation settings, both in simulation and in the real world. Our experiments seek to address the following questions:

- Do contrastive learning methods learn better latent spaces and forward models for planning in deformable object manipulation tasks?
- What aspects of our contrastive learning methods contribute the most to performance?
- Can we successfully manipulate deformable objects on a real-world robot?

### A. Environments and Tasks

To simulate deformable objects such as cloth and rope, we used the Deep Mind Control [51] platform with MuJoCo 2.0 [53]. We use an overhead camera that renders  $64 \times 64 \times 3$  RGB images as input observations for training our method.

We design the following tasks in simulation:

**1. Rope:** The rope is represented by 25 geoms in simulation with a four-dimensional action space: the first 2 are the pixel pick point on the rope, and the last 2 are the  $x, y$  delta direction to perturb the rope. At the start of each episode, the rope's state is randomly initialized by applying 120 random actions.

**2. Cloth:** The cloth is represented by a  $9 \times 9$  grid of geoms in simulation with a five-dimensional action space: the first 2 are the pixel pick point on the cloth, and the last 3 are the  $x, y, z$  delta direction to perturb the cloth. At the start of each episode, the cloth's state is randomly initialized by applying 50 random actions. In MuJoCo 2.0, the skin of the cloth can be changed by using images taken of a real cloth.

For both rope and cloth environments, we evaluate our method by planning to a desired goal state image and computing the sum of the pairwise geom distances between the achieved and true goal states. We observe that taking an average of 1000 trials suffices to maintain high-confidence evaluation estimates.

### B. Data Collection

Since collecting real-world data on robots is expensive, our method seeks to address this problem by collecting randomly perturbed rope and cloth data in simulation. Using random perturbations allows for a diverse set of deformable objects

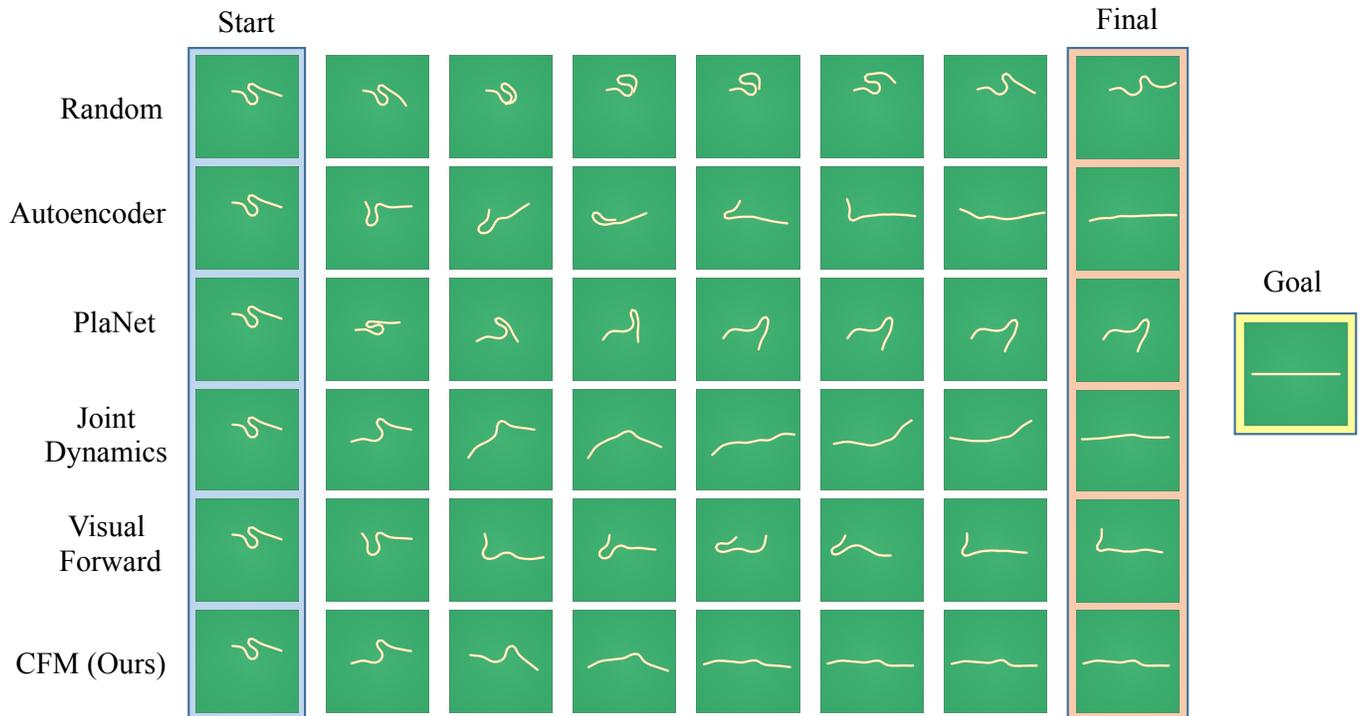


Fig. 3: Trajectories for each of the baselines within the simulator, all starting from the same start state and having the same end goal of a horizontal line. Each trajectory was run for 20 actions. Note that our method (CFM) reaches the goal state significantly faster than the baselines.

TABLE I: Quantitative comparisons between different model-based learning methods on rope and cloth manipulation tasks. The metric is the sum of pairwise geom distances between the final observation and goal state, where lower distance is more accurate.

			Rope			Cloth	
	Horizontal	Vertical	45°	135°	Random	Flat	Random
<b>Random Policy</b>	4.75	4.93	4.80	4.87	5.73	7.98	10.12
<b>Autoencoder</b>	1.72 ± 0.31	3.24 ± 1.28	<b>2.11 ± 0.51</b>	2.49 ± 0.64	4.308 ± 1.16	3.24 ± 0.29	4.82 ± 0.0
<b>PlaNet</b>	1.81 ± 0.13	3.36 ± 0.78	2.31 ± 0.72	2.38 ± 0.20	3.037 ± 0.24	4.12 ± 0.21	5.06 ± 0.02
<b>Joint Dynamics Model</b>	2.13 ± 0.66	4.33 ± 0.85	3.88 ± 0.95	4.02 ± 0.85	1.78 ± 0.09	4.24 ± 0.06	4.70 ± 0.03
<b>Visual Forward Model</b>	2.09 ± 0.13	<b>2.65 ± 0.27</b>	2.55 ± 0.34	2.27 ± 0.17	4.77 ± 0.18	<b>2.20 ± 0.05</b>	4.65 ± 0.10
<b>CFM (Ours)</b>	<b>0.58 ± 0.09</b>	3.08 ± 1.19	2.29 ± 1.42	<b>2.24 ± 0.90</b>	<b>1.52 ± 0.10</b>	2.69 ± 0.25	<b>3.97 ± 0.16</b>
			Rope (With DR)			Cloth (With DR)	
	Horizontal	Vertical	45°	135°	Random	Flat	Random
<b>Random Policy</b>	4.75	4.93	4.80	4.87	5.73	7.975	10.12
<b>Autoencoder</b>	3.29 ± 1.08	3.70 ± 1.47	3.19 ± 1.14	3.30 ± 1.14	4.31 ± 1.16	6.26 ± 1.23	7.08 ± 2.22
<b>PlaNet</b>	2.35 ± 0.56	4.06 ± 1.84	3.73 ± 1.66	3.58 ± 1.46	3.04 ± 0.24	8.74 ± 0.55	10.10 ± 1.56
<b>Joint Dynamics Model</b>	1.01 ± 0.40	2.29 ± 0.10	1.35 ± 0.59	1.82 ± 0.50	1.78 ± 0.09	4.17 ± 0.17	4.64 ± 0.20
<b>Visual Forward Model</b>	3.05 ± 0.45	5.65 ± 0.37	5.37 ± 0.90	5.11 ± 1.04	4.77 ± 0.18	6.64 ± 0.66	6.07 ± 0.52
<b>CFM (Ours)</b>	<b>0.88 ± 0.21</b>	<b>1.20 ± 0.07</b>	<b>0.99 ± 0.07</b>	<b>0.99 ± 0.17</b>	<b>1.38 ± 0.03</b>	<b>3.99 ± 0.15</b>	<b>4.40 ± 0.06</b>

and interactions for learning the latent space and dynamics model. We collect 4000 trajectories of length 50 for rope (200k samples), and 8000 trajectories of length 50 for cloth (400k samples).

### C. Baselines

To show the substantial improvements of our model over prior methods, we compare our method against several baselines: a random policy, a visual forward model, an autoencoder trained jointly with a latent dynamics model, PlaNet [14], and a joint dynamics model [1]. In order to ensure that pick

points are always on the rope or cloth, we constrain our pick points using a binary segmentation of the observation image computed by RGB thresholding. During planning, all methods use MPC with one-step prediction.

- **Random Policy:** We sample pick actions uniformly over the binary segmentation, and place actions are sampled uniformly random in a unit square centered around the pick location.
- **Visual Forward Model:** We train a forward model similar to Kaiser et al. [19] to perform modeling and planning purely through pixel space.

- **Autoencoder:** We learn a simple latent space model by jointly training a classical autoencoder with a forward dynamics model. The autoencoder learns to minimize the  $\ell_2$ -distance between reconstructed and actual images [24].
- **PlaNet:** We train PlaNet [14], a stochastic variant of an autoencoder, as another latent space model. PlaNet models a sequential VAE and optimizes a temporal variational lower bound.
- **Joint Dynamics Model:** We jointly learn a forward and inverse model following Agrawal et al. [1].

For consistency across all latent space models, we use a latent size of 8 for both the rope and the cloth environments. For all methods, we sample 100 possible one-step trajectories when performing closed-loop planning. See Figure 3 for example trajectories from each baseline in comparison to our method.

#### D. Training Details

We used the same encoder architectures for all models. The encoder architecture is a series of 6 2D convolutions of kernel sizes [3, 4, 3, 4, 4, 4], strides [1, 2, 1, 2, 2, 2], and filter sizes [64, 64, 64, 128, 256, 256] respectively. We add Leaky ReLU [27] activation in between each convolutional layer. Finally, the output is flattened and fed into a fully connected layer to produce the latent  $z$ . The forward model is a multi-layer perceptron (MLP) with two hidden layers of size 32 which outputs the parameters for a linear transformation on  $z_t$ . Specifically for our method (CFM), we use the other batch elements as our negative samples for a total of 127 negative samples per positive pair. For PlaNet, following Hafner et al. [14], the decoder architecture is a dense layer followed by 4 transposed convolutions with kernel size 4 and stride 2 to upscale to the size of the  $64 \times 64$  image. The visual forward models follow the same convolutional encoder and decoder architectures as the previous model, with action conditioning implemented in a similar way to Kaiser et al. [19] where actions are processed by a separate dense layer at each resolution, multiplied channel-wise, and broadcasted spatially. The images and actions were centered and scaled to the range of  $[-1, 1]$ . We trained all models with batch size 128, learning rate  $10^{-3}$ , and an Adam optimizer [21] for 30 epochs. Each model was trained on a single NVIDIA TitanX GPU, taking roughly 1–2 hours. All of our simulated environments, evaluation metrics, and training code will be publicly released.

#### E. Does Using Contrastive Models Improve Performance?

In this section, we compare the results of using our method with those of our baselines, analyzing the advantages and benefits that contrastive models bring over prior methods. Consider a naive baseline where we replace the InfoNCE loss with an MSE loss. This is equivalent to jointly fitting an encoder and dynamics model that minimizes  $\|f_\phi(g_\theta(o_t), a_t) - z_{t+1}\|^2$ . We can see that the optimal solution is for the encoder to encode all observations to a constant vector to achieve zero loss. To prevent this form of a degenerate solution, we are required to regularize our latent space in some way. Both

prior methods and contrastive learning do this in different ways so we analyzed which methods performed better over others. Table I shows the quantitative results comparing our method against baselines in different rope and cloth environments, with and without domain randomization for robot transfer. Note that our method does better on all randomly sampled goals with and without domain randomization, indicating stronger generalization in latent spaces for planning. Figure 3 shows example simulator trajectories for each baseline. Each trajectory has the same starting location, same goal image, and was run for 20 actions.

An autoencoder regularizes its latent space by requiring additionally training a decoder to learn to reconstruct  $o_t$  from  $z_t$ . The model does well in some scenarios, such as a  $45^\circ$  diagonal, but performs poorly when domain randomization is introduced to allow for transfer to a real robot. This is most likely because the autoencoder is optimized to have pixel-level perfect reconstructions, so features such as lighting and color must be encoded in the latent space even when they are not needed for the task. PlaNet behaves similarly to the autoencoder, as it is also a form of a stochastic autoencoder. It performs reasonably competitive with our method but again fails when domain randomization is introduced.

The joint dynamics model regularizes its latent space by jointly learning an inverse model with the forward model. The joint model performed the best across all the baselines when moving to domain randomized data. However, our method still outperforms the joint model for every task.

The visual forward model is the only method that plans in pixel space. It generally performs poorly for tasks with objects with low area coverage, such as the different rope goal orientations, but does better than our method on the cloth flattening task. However, since the forward model operates purely in pixel space, it unsurprisingly suffers from a sharp degradation in performance when introducing domain randomization. As such, it generalizes poorly to the real robot setting.

TABLE II: Ablation experiments on forward model architecture and similarity functions, using the same evaluation metric as Table I (lower is better).

	Rope	Cloth
<b>Linear</b>	$2.30 \pm 0.31$	$4.10 \pm 0.03$
<b>MLP</b>	$1.56 \pm 0.10$	$3.95 \pm 0.10$
<b>Log-bilinear Similarity</b>	$3.25 \pm 0.43$	$4.16 \pm 0.15$
<b>Ours</b>	$1.40 \pm 0.02$	$3.81 \pm 0.09$

#### F. Ablations on Contrastive Models

In this section, we perform an ablation study on our method, examining the impact of architectural designs on performance. We ablate over two aspects of our method: the forward model architecture, and the contrastive similarity function. For the forward model, our method uses a Multi-Layer Perceptron (MLP) that outputs the parameters of a linear function that is then applied to  $z_t$ . For the contrastive similarity function, our method follows Equation 2. The quantitative results, measured

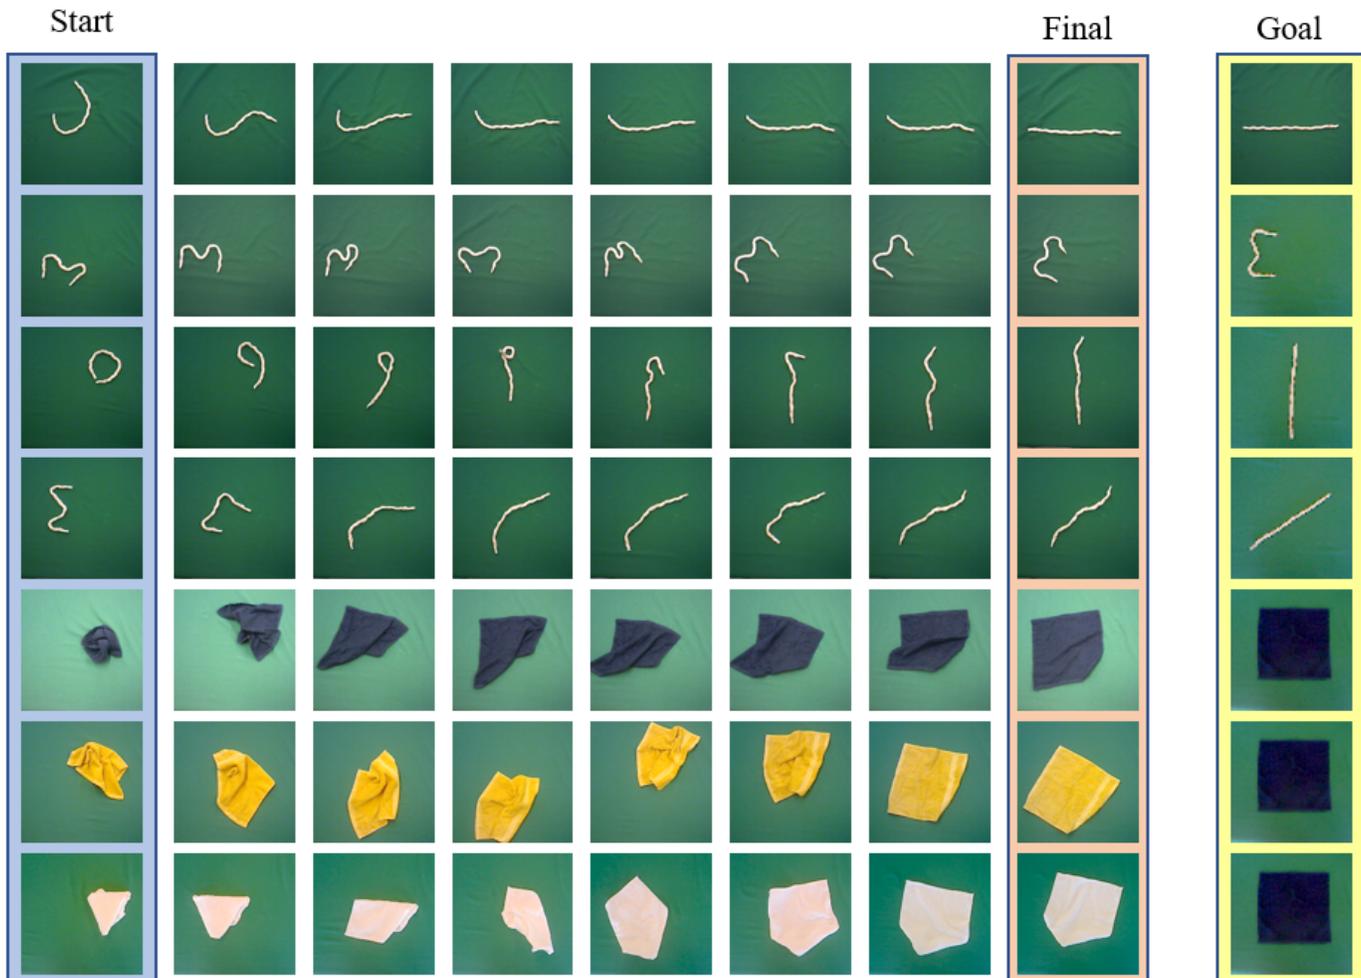


Fig. 4: Each row represents one trajectory using our contrastive forward policy. The rope task uses 40 actions between the start and final states, while the cloth task uses 100 actions.

TABLE III: The maximum intersection area in pixels between the goal image and observation images averaged over all seeds

Robot Experiments (Intersection in pixels)	Rope (Horizontal)	Rope (Vertical)	Rope (45°)	Rope (135°)	Rope (Squiggle)	Cloth (Flat)
Random Policy	6.880	14.727	13.662	4.266	0.049	462.513
Autoencoder	5.526	3.334	3.862	7.499	3.419	603.927
Joint Dynamic Model	17.722	23.636	33.631	21.267	18.311	772.303
<b>Contrastive Forward Model (Ours)</b>	<b>32.827</b>	<b>36.387</b>	<b>33.891</b>	<b>38.952</b>	<b>20.711</b>	<b>1001.082</b>

as the sum of pairwise geom distances between the final and goal images, appear in Table II.

1) *Contrastive Similarity Functions*: We compare using our similar function with the original InfoNCE similarity function in Oord et al. [37], the log-bilinear similarity function  $h(z_1, z_2) = \exp(z_1^T z_2)$ . We achieve the largest boost in performance when switching to our similarity function, as it is more in line with the minimization objective of learning a correct forward model, whereas the log-bilinear model only encourages alignment (as opposed to closeness) of embedding vectors.

2) *Forward models architectures*: We experiment with a few different forward model architectures: linear, a small MLP, and a small MLP that outputs parameters for a linear transformation. As expected, the biggest drop in performance occurs when learning the simpler linear dynamics model, and a slight drop when using an MLP for both rope and cloth tasks. This demonstrates the need for more complex models for latent forward-dynamics learning.

#### G. Real Robot Experiments

1) *Real Robot Setup*: We use a PR2 robot to perform our experiments and an overhead camera looking down on the deformable objects to get the RGB image inputs. To ensure the

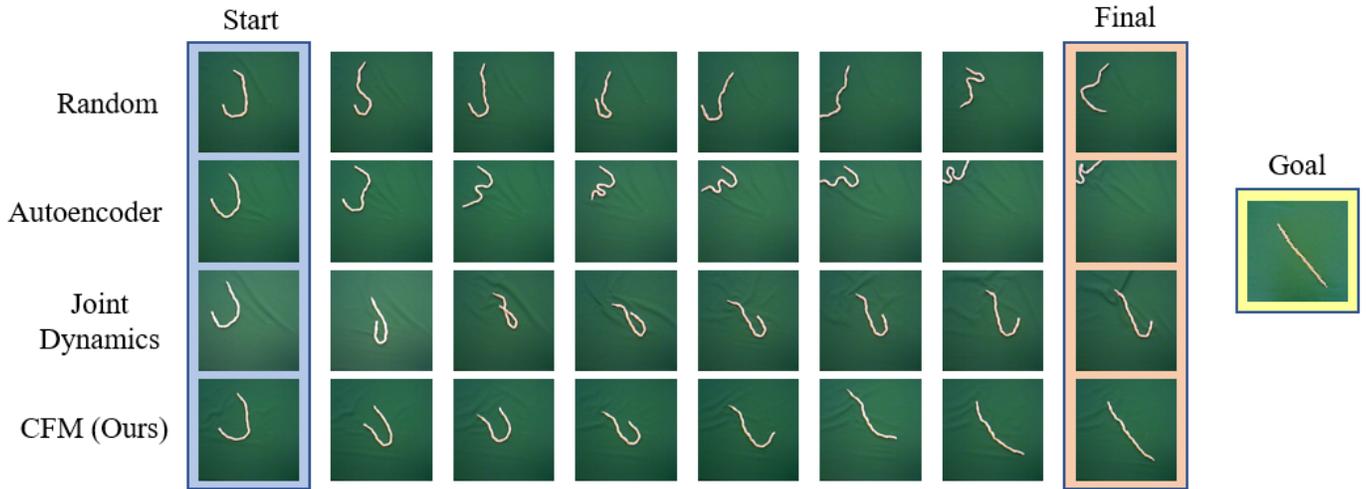


Fig. 5: Rope trajectories for each baseline and our method applied on a real robot, all with the same start state, and  $135^\circ$  goal rope orientation.

policy learned in the simulator transfers over to the real world, we apply domain randomization by changing the lighting, texture, friction, damping, inertia, and mass of the object during every training step within the simulator. We also use a pick and place strategy to mimic the same four-dimensional actions within the simulator.

To compute the actions, we employ a model predictive control (MPC) approach of replanning our action at each time step based on the previous image. We segment the rope/cloth against the background to get the list of valid pick locations of the object. We then generate possible actions by uniformly sampling 100 random deltas in  $[-1, 1]$  combined with randomly chosen start locations. We feed these into our forward model along with the encoding of our start image to get the latent encoding for each of the next prospective states. To pick the optimal action, we find the location and delta that minimizes the Euclidean distance from these next states to our goal state and return this action to the robot. The delta from the policy is on the scale of  $[-1, 1]$  for both  $x$  and  $y$  coordinates, and we rescale this to  $[-5, 5]$  pixels. On the robot side, we use a learned linear mapping to transform from the image’s pixel values to Cartesian coordinates that the robot uses. To emulate the simulator, the robot’s left arm motion is to go to the start location, go down and close the gripper, move up, move to the new location, move down and open the gripper, where the height of the gripper is hard-coded to some manually tuned value.

2) *Evaluation Metrics*: We use three baselines along with our contrastive method for real-world evaluation. The first is random actions and the others are the two policies that performed the best with domain randomization: the autoencoder and the joint dynamics model [1]. For the rope, all the models are evaluated on five goal states: horizontal, vertical, straight line at  $45^\circ$ , straight line at  $135^\circ$ , and a squiggly rope on left. For the cloth, the models are evaluated on one goal state, a flat blue cloth with no rotation. The metric we use is the intersection in pixels between the segmented

final image and the segmented goal image. We prefer this instead of intersection over union (IOU) since the objects have the same shape so the union normalization is unnecessary. Additionally, the simpler intersection values provide more insight for comparisons than IOU. The models are run for 40 actions on the rope or 100 actions on the cloth, and the image after each action is stored as an observation. Among all the observations, the one with the highest intersection with the goal is chosen for each method. To account for different seeds, we use 4 starting locations for our contrastive method and 2 starting locations for the baselines, with the scores being averaged across the different start locations. For the cloth, the seed also involves different colors of cloths (blue, gold, white).

The specific evaluation metrics are found in Table III which shows that our model performed the best for all the rope and cloth tasks. The joint dynamics model is the second best and got close results to ours on the  $45^\circ$  and squiggly rope tasks. Some example trajectories from our model are seen from a forward view in Figure 1 and from an overhead view in Figure 4. Visual comparisons between our method and baseline methods on the real robot are found in Figure 5. We see that our method more accurately plans towards correct goal states compared to the baselines.

## V. CONCLUSION

In this paper, we propose a contrastive learning approach for predictive modeling of deformable objects. We show that contrastive learning learns stronger and more plannable latent representations compared to existing methods. Since our method only requires collecting random data in an environment, it allows for easier transfer to real robots without the need for real-world training.

## ACKNOWLEDGEMENTS

We thank AWS for computing resources. We also gratefully acknowledge the support from Berkeley DeepDrive, NSF, and the ONR Pecase award.

## REFERENCES

- [1] Pulkit Agrawal, Ashvin Nair, Pieter Abbeel, Jitendra Malik, and Sergey Levine. Learning to poke by poking: Experiential learning of intuitive physics. *NIPS*, 2016.
- [2] Marcin Andrychowicz, Bowen Baker, Maciek Chociej, Rafal Jozefowicz, Bob McGrew, Jakub Pachocki, Arthur Petron, Matthias Plappert, Glenn Powell, Alex Ray, et al. Learning dexterous in-hand manipulation. *arXiv preprint*, 2018.
- [3] Dmitry Berenson. Manipulation of deformable objects without modeling and simulating deformation. In *IROS*, 2013.
- [4] Eduardo F Camacho and Carlos Bordons Alba. *Model predictive control*. Springer Science & Business Media, 2013.
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint*, 2020.
- [6] Xi Chen, Yan Duan, Rein Houthoof, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *NIPS*, 2016.
- [7] Yan Duan, Xi Chen, Rein Houthoof, John Schulman, and Pieter Abbeel. Benchmarking deep reinforcement learning for continuous control. In *ICML*, 2016.
- [8] Frederik Ebert, Chelsea Finn, Sudeep Dasari, Annie Xie, Alex Lee, and Sergey Levine. Visual foresight: Model-based deep reinforcement learning for vision-based robotic control. *arXiv preprint arXiv:1812.00568*, 2018.
- [9] Nabil Essahbi, Belhassen Chedli Bouzgarrou, and Grigore Gogu. Soft material modeling for robotic manipulation. In *Applied Mechanics and Materials*, 2012.
- [10] Chelsea Finn and Sergey Levine. Deep visual foresight for planning robot motion. In *ICRA*, 2017.
- [11] Barbara Frank, Cyrill Stachniss, Nichola Abdo, and Wolfram Burgard. Efficient motion planning for manipulation robots in environments with deformable objects. In *IROS*, 2011.
- [12] Abhinav Gupta, Adithyavairavan Murali, Dhiraj Prakashchand Gandhi, and Lerrel Pinto. Robot learning in homes: Improving generalization and reducing dataset bias. In *NeurIPS*, 2018.
- [13] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, Abhishek Gupta, Pieter Abbeel, et al. Soft actor-critic algorithms and applications. *arXiv preprint*, 2018.
- [14] Danijar Hafner, Timothy Lillicrap, Ian Fischer, Ruben Villegas, David Ha, Honglak Lee, and James Davidson. Learning latent dynamics for planning from pixels. *arXiv preprint*, 2018.
- [15] Dominik Henrich and Heinz Wörn. *Robot manipulation of deformable objects*. Springer Science & Business Media, 2012.
- [16] Zhe Hu, Peigen Sun, and Jia Pan. Three-dimensional deformable object manipulation using fast online gaussian process regression. *RAL*, 2018.
- [17] Biao Jia, Zhe Hu, Zherong Pan, Dinesh Manocha, and Jia Pan. Learning-based feedback controller for deformable object manipulation. *arXiv preprint*, 2018.
- [18] P Jiménez. Survey on model-based manipulation planning of deformable objects. *Robotics and computer-integrated manufacturing*, 2012.
- [19] Lukasz Kaiser, Mohammad Babaeizadeh, Piotr Milos, Blazej Osinski, Roy H Campbell, Konrad Czechowski, Dumitru Erhan, Chelsea Finn, Piotr Kozakowski, Sergey Levine, et al. Model-based reinforcement learning for atari. *arXiv preprint*, 2019.
- [20] Fouad F Khalil and Pierre Payeur. Dexterous robotic manipulation of deformable objects with multi-sensory feedback-a review. In *Robot Manipulators Trends and Development*. 2010.
- [21] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [22] Vikash Kumar, Emanuel Todorov, and Sergey Levine. Optimal control with learned local models: Application to dexterous manipulation. In *ICRA*, 2016.
- [23] Thanard Kurutach, Aviv Tamar, Ge Yang, Stuart J Russell, and Pieter Abbeel. Learning plannable representations with causal infogan. In *NeurIPS*, 2018.
- [24] Sascha Lange and Martin Riedmiller. Deep auto-encoder neural networks in reinforcement learning. In *IJCNN*, 2010.
- [25] Sergey Levine, Peter Pastor, Alex Krizhevsky, and Deirdre Quillen. Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *ISER*, 2016.
- [26] Timothy P Lillicrap, Jonathan J Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. *arXiv preprint*, 2015.
- [27] Andrew L Maas, Awni Y Hannun, and Andrew Y Ng. Rectifier nonlinearities improve neural network acoustic models. In *ICML*, 2013.
- [28] J. Mahler, F. T. Pokorny, B. Hou, M. Roderick, M. Laskey, M. Aubry, K. Kohlhoff, T. Kröger, J. Kuffner, and K. Goldberg. Dex-net 1.0: A cloud-based network of 3d objects for robust grasp planning using a multi-armed bandit model with correlated rewards. In *ICRA*, 2016.
- [29] Jeremy Maitin-Shepard, Marco Cusumano-Towner, Jinna Lei, and Pieter Abbeel. Cloth grasp point detection based on multiple-view geometric cues with application to robotic towel folding. In *ICRA*, 2010.
- [30] Jan Matas, Stephen James, and Andrew J Davison. Sim-to-real reinforcement learning for deformable object manipulation. *arXiv preprint*, 2018.
- [31] Dale McConachie and Dmitry Berenson. Estimating model utility for deformable object manipulation using multiarmed bandit methods. *IEEE Transactions on Automation Science and Engineering*, 2018.
- [32] Dale McConachie, Mengyao Ruan, and Dmitry Berenson. Interleaving planning and control for deformable object manipulation. In *ISRR*, 2017.
- [33] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *NIPS*, 2013.
- [34] Mark Moll and Lydia E Kavraki. Path planning for deformable linear objects. *T-RO*, 2006.
- [35] Anusha Nagabandi, Gregory Kahn, Ronald S Fearing, and Sergey Levine. Neural network dynamics for model-based deep reinforcement learning with model-free fine-tuning. In *ICRA*, 2018.
- [36] Ashvin Nair, Dian Chen, Pulkit Agrawal, Phillip Isola, Pieter Abbeel, Jitendra Malik, and Sergey Levine. Combining self-supervised learning and imitation for vision-based rope manipulation. In *ICRA*, 2017.
- [37] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint*, 2018.
- [38] Piotr Pierański, Sylwester Przybył, and Andrzej Stasiak. Tight open knots. *The European Physical Journal E*, 2001.
- [39] Lerrel Pinto and Abhinav Gupta. Supersizing self-supervision: Learning to grasp from 50k tries and 700 robot hours. *ICRA*, 2016.
- [40] Samuel Rodriguez, Xinyu Tang, Jyh-Ming Lien, and Nancy M Amato. An obstacle-based rapidly-exploring random tree. In *ICRA*, 2006.
- [41] Mitul Saha and Pekka Isto. Manipulation planning for deformable linear objects. *T-RO*, 2007.
- [42] John Schulman, Jonathan Ho, Cameron Lee, and Pieter Abbeel. Generalization in robotic manipulation through the use of non-rigid registration. In *ISRR*, 2013.
- [43] John Schulman, Alex Lee, Jonathan Ho, and Pieter Abbeel. Tracking deformable objects with point clouds. In *ICRA*, 2013.
- [44] John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan,

- and Philipp Moritz. Trust region policy optimization. In *ICML*, 2015.
- [45] Daniel Seita, Nawid Jamali, Michael Laskey, Ajay Kumar Tanwani, Ron Berenstein, Prakash Baskaran, Soshi Iba, John Canny, and Ken Goldberg. Deep transfer learning of pick points on fabric for robot bed-making. *arXiv preprint*, 2018.
- [46] Daniel Seita, Aditya Ganapathi, Ryan Hoque, Minh Hwang, Edward Cen, Ajay Kumar Tanwani, Ashwin Balakrishna, Brijen Thananjeyan, Jeffrey Ichnowski, Nawid Jamali, Katsu Yamane, Soshi Iba, John Canny, and Ken Goldberg. Deep imitation learning of sequential fabric smoothing policies. *arXiv preprint*, 2019.
- [47] Karun B Shimoga. Robot grasp synthesis algorithms: A survey. *IJRR*, 1996.
- [48] Jerzy Smolen and Alexandru Patriciu. Deformation planning for robotic soft tissue manipulation. In *Advances in Computer-Human Interactions*, 2009.
- [49] Akash Srivastava, Lazar Valkov, Chris Russell, Michael U Gutmann, and Charles Sutton. Veegan: Reducing mode collapse in gans using implicit variational learning. In *NIPS*, 2017.
- [50] Jan Stria, Daniel Prusa, Vaclav Hlavac, Libor Wagner, Vladimir Petrik, Pavel Krsek, and Vladimir Smutny. Garment perception and its folding using a dual-arm robot. In *IROS*, 2014.
- [51] Yuval Tassa, Yotam Doron, Alistair Muldal, Tom Erez, Yazhe Li, Diego de Las Casas, David Budden, Abbas Abdolmaleki, Josh Merel, Andrew Lefrancq, et al. Deepmind control suite. *arXiv preprint*, 2018.
- [52] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *arXiv preprint*, 2019.
- [53] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *IROS*, 2012.
- [54] Takahiro Wada, Shinichi Hirai, Sadao Kawamura, and Norimasa Kamiji. Robust manipulation of deformable objects by a simple pid feedback. In *ICRA*, 2001.
- [55] Hidefumi Wakamatsu, Eiji Arai, and Shinichi Hirai. Knotting/unknotted manipulation of deformable linear objects. *IJRR*, 2006.
- [56] Angelina Wang, Thanard Kurutach, Kara Liu, Pieter Abbeel, and Aviv Tamar. Learning robotic manipulation through visual planning and acting. *RSS*, 2019.
- [57] Grady Williams, Nolan Wagener, Brian Goldfain, Paul Drews, James M Rehg, Byron Boots, and Evangelos A Theodorou. Information theoretic mpc for model-based reinforcement learning. In *ICRA*, 2017.
- [58] Yilin Wu, Wilson Yan, Thanard Kurutach, Lerrel Pinto, and Pieter Abbeel. Learning to manipulate deformable objects without demonstrations. *arXiv preprint*, 2019.
- [59] Hanna Yousef, Mehdi Boukallel, and Kaspar Althoefer. Tactile sensing for dexterous in-hand manipulation in robotics—a review. *Sensors and Actuators A: physical*, 2011.