

# TraKDis: A Transformer-based Knowledge Distillation Approach for Visual Reinforcement Learning with Application to Cloth Manipulation

Wei Chen<sup>1</sup> *Graduate Student Member, IEEE* and Nicolas Rojas<sup>2</sup> *Member, IEEE*

**Abstract**—Approaching robotic cloth manipulation using reinforcement learning based on visual feedback is appealing as robot perception and control can be learned simultaneously. However, major challenges result due to the intricate dynamics of cloth and the high dimensionality of the corresponding states, what shadows the practicality of the idea. To tackle these issues, we propose *TraKDis*, a novel Transformer-based Knowledge Distillation approach that decomposes the visual reinforcement learning problem into two distinct stages. In the first stage, a privileged agent is trained, which possesses complete knowledge of the cloth state information. This privileged agent acts as a teacher, providing valuable guidance and training signals for subsequent stages. The second stage involves a knowledge distillation procedure, where the knowledge acquired by the privileged agent is transferred to a vision-based agent by leveraging pre-trained state estimation and weight initialization. *TraKDis* demonstrates better performance when compared to state-of-the-art RL techniques, showing a higher performance of 21.9%, 13.8%, and 8.3% in cloth folding tasks in simulation. Furthermore, to validate robustness, we evaluate the agent in a noisy environment; the results indicate its ability to handle and adapt to environmental uncertainties effectively. Real robot experiments are also conducted to showcase the efficiency of our method in real-world scenarios.

## I. INTRODUCTION

Manipulating deformable objects, especially cloth and other fabrics, with robotic technology is a considerable challenge that has far-reaching and widespread applications in domestic, medical, and industrial settings. Important tasks such as assistive dressing, fabric manufacturing automation, and household cleaning tasks all require robust, delicate, and precise manipulation of cloth to achieve their goals [1]. While the manipulation of rigid-body objects is relatively well-developed, that of cloth remains in its infancy.

Reinforcement learning (RL) is a popular control tool for robotic manipulation tasks that has been utilised in cloth manipulation successfully [2]–[5]. However, RL agents are typically trained using cloth state information, such as the exact particle positions of keypoints on the cloth. This allows the agent to achieve satisfactory performance in simulated environments; however, obtaining such state information in the real world is particularly challenging. Rather than relying on exact object states, some works infer them from RGB image data in a visual manipulation policy [3], [6]. However, while this is often successful in rigid-object manipulation [7],

[8], cloth and fabric exhibit high self-occlusion and are often lacking in tracking observable features such as edges and corners, making visual manipulation difficult. Further still, visual data such as RGB image data always presents higher dimensionality than state information data, which makes directly training a visual RL policy more challenging.

Most of the existing visual RL-based works for cloth manipulation employ an asymmetric actor-critic architecture for their policy training. This is implemented by using high-dimension observation (RGB images) and low-dimension environment states as the input for the actor and critic, respectively [9]. Although this asymmetric architecture has been applied in rigid-body and deformable object manipulation [3], [6], [10], the utilization of actor-critic also comes with the limitation that all decisions are made based on the Markov assumption, which may cause robots to lose historical information [11], [12].

In a highly dynamic environment such as cloth manipulation, the current state of cloth may depend not only on the last state but also on a sequence of past states. The loss of historical information may influence the controllability and observability of the overall control agent. To address this issue, in contrast to previous works [3], [6], we adopt a transformer-based RL approach for the learning framework. We argue that the adoption of the transformer model yields twofold benefits. First, the transformer architecture can process information with a large capacity and a long memory, which holds more significant promise for accurately predicting and understanding complex trajectory sequences [13]. Second, transformers excel at capturing contextual relationships, which enhance knowledge transfer by understanding data nuances and dependencies. Pre-training on a large dataset enables recognition and comprehension of contextual patterns, aiding fine-tuning specific tasks [14].

This work then proposes *TraKDis*, a novel **T**ransformer-based **K**nowledge **D**istillation (KD) approach for learning visual cloth manipulation. Our goal here is to solve the highly dynamic cloth manipulation tasks with only RGB image input. Unlike the asymmetric input of RGB images and states in actor-critic approaches [6], we apply KD for the learning of vision-based tasks. While most KD procedures focus on compressing multiple teachers' knowledge into one student model, *TraKDis* can be considered as a special *one-to-one* case [15]. As shown in Fig.1, *TraKDis* decomposes the learning of the vision agent into two stages. First, we train a privileged agent that takes the privileged cloth state information as input. A novel KD procedure is then conducted to distil the knowledge from the privileged agent to

This work was supported in part by the China Scholarship Council and the Dyson School of Design Engineering, Imperial College London.

<sup>1</sup>Wei Chen is with Dyson School of Design Engineering, Imperial College London, 25 Exhibition Road, London, SW7 2DB, UK. w.chen21@imperial.ac.uk N. Rojas<sup>2</sup> is with The AI Institute, Cambridge, MA, USA.

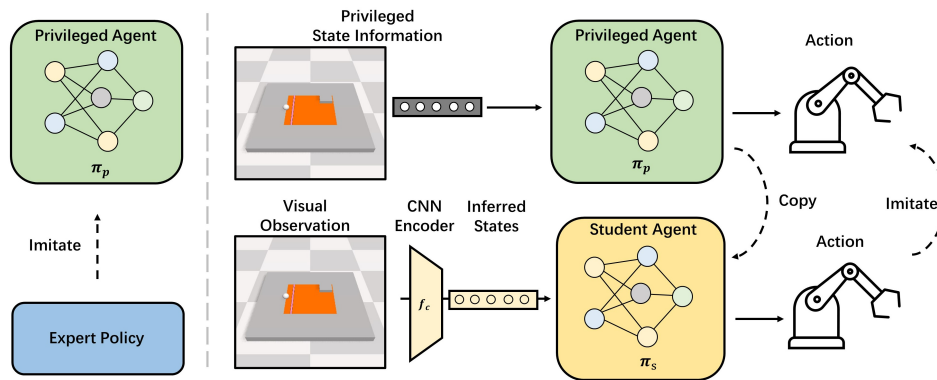


Fig. 1. Pipeline of *TraKDis* for training visual RL agents for cloth manipulation. A privilege agent is trained to imitate the expert policy with privileged state information (cloth particle locations) to obtain a robust performance. Then, a student agent that can only access partial observation (RGB images) is trained to imitate the privileged agent. By leveraging a CNN encoder and weight copy, the student policy can realize an enhanced performance with limited observation information.

the vision-based (student) agent. Our proposed KD approach consists of two major components: a state estimation encoder and the weight transfer. Specifically, the pre-trained CNN state estimation encoder is applied as the *prior* to project the high-dimensional images into estimated low-dimensional states. This aims to reduce the domain gap between visual observation and states during the distillation. Since the two models share the same architecture, we also initialize the student policy by weight transfer from the privilege agent at the beginning of distillation, which shows significant improvement in facilitating knowledge distillation.

We highlight our key contribution as follows: (1) *TraKDis*: a novel transformer-based knowledge distillation framework for learning visual cloth manipulation tasks. (2) We propose a novel approach for the knowledge distillation procedure by leveraging a state estimation encoder and pre-trained weight. The student model (image) can leverage the teacher model (states) knowledge, significantly increasing the model’s performance and training efficiency. (3) With the proposed learning framework, *TraKDis* outperforms other state-of-the-art (SOTA) baselines in several comparison experiments.

## II. RELATED WORKS

### A. Cloth Manipulation by Data-driven Methods

Due to the characteristics of high elasticity, low controllability, and complex dynamics, cloth manipulation remains a challenge in the field of robotic manipulation [16], [17]. In recent years, some data-driven methods approaches have gained popularity. Previous works tried to use deep learning (DL)-based methods to detect key cloth regions for downstream tasks such as bed-making, cloth grasping, cloth flattening and garment hanging [18]–[23].

By learning from trial and error, RL and RL combined with Imitation Learning (IL) methods have made significant progress in deformable object manipulation, especially cloth manipulation tasks. Using the combination of deep deterministic policy gradient (DDPG), Hindsight Experience Replay (HER) and expert demonstration, [4] tries to learn the dynamic cloth folding tasks. However, one preliminary

of this approach is the access to the state information of cloth during the whole manipulation procedure. In the real world, however, extracting such ground truth state information is not always available for a robot because of the large configuration space of deformable objects. On the other hand, rather than state information, it is more feasible for a robot to use high-dimensional observations such as RGB images as inputs to its control policy [3].

Due to the deformable nature of cloth and the typically high-dimensional configuration parameters, it is extremely challenging to infer the cloth states from only visual input. This makes vision-based cloth manipulation more difficult. With some recent progress in RL, CURL [24] and PlaNet [25] are proposed for the training of image-based agents. Although several image-based RL methods, including CURL and PlaNet have been used as the benchmark to test the performance of visual cloth manipulation tasks in [3], the results are far below the optimal performance (state-based agent) on many tasks. Furthermore, by combining the expert demonstration and RL, [6] has adopted an asymmetric architecture for Learning from Demonstration (LfD) methods, where the actor acts based solely on environment images while the critics learn from both image and state information. However, this still results in a significant gap when the agent performs tasks with only image observation compared to state input.

### B. Transformer

Benefiting from the attention mechanism and its significant advantages of sequence modelling, transformers are initially applied in the field of Natural Language Processing (NLP) [26]. Being attributed to its high-quality global contextual learning and friendly parallel computation, transformer architecture has emerged as a high-capacity learning framework over traditional sequence prediction methods such as Recurrent Neural Networks (RNNs) or Long Short-Term Memory Networks (LSTMs).

With some breakthroughs in transformers, the application area of transformers also extends to vision and reinforcement learning (RL) [27]–[29]. For instance, [27] proposed the




Fig. 2. Model Training of the privileged agent. We first conduct offline training using the expert data collected from the human-designed heuristic algorithm. An online fine-tuning is then applied for the trained model to obtain a better performance

vision transformer (ViT) which shows better accuracy over the traditional CNN with a large training dataset. Furthermore, researchers have also utilized transformer architecture in decision-making agents, proposing the decision transformer (DT) model, which re-frames the sequence modelling problem as an action prediction problem based on a series of historical states and reward-to-go [28]. Training in a supervised fashion, DT tries to minimize the error between predicted and ground-truth actions by inputting past state and reward-to-go sequences. Some recent work has shown DT has superiority over many state-of-the-art offline RL algorithms [28]. Furthermore, more variants of DT have been proposed to extend the DT model to the online training scenario [28] and multi-agent scenario [30].

### C. Knowledge Distillation

Knowledge distillation is a transfer learning (TL) approach that aims to transfer knowledge from the teacher model to the student model [15]. The original idea of knowledge distillation (KD) is to transfer knowledge from a large and complex model (the teacher) to a small and simplified model (the student) [31]. This transfer is usually achieved by minimizing the Kullback–Leibler (KL) divergence of teacher and student outputs.

Some recent work has integrated state estimation and TL in robotics for locomotion control of legged robots. By integrating a Temporal-Convolutional Neural Network (TCNN) or temporal-MLP, temporal information can be applied to enhance state estimation for the learning of the student policy in these methods [32], [33]. In contrast, *TraKDis* utilizes the DT architecture for temporal decision-making, using the state input estimated by the CNN encoder.

## III. METHODS

### A. Preliminary

We formulate the cloth manipulation environment as a partially observable Markov decision process (POMDP). The tuple of  $\langle \mathcal{A}, \mathcal{S}, \mathcal{O}, \mathcal{P}, \mathcal{R} \rangle$  are used to describe the POMDP, which consists of actions  $a \in \mathcal{A}$ , states  $s \in \mathcal{S}$ , observations  $o \in \mathcal{O}$ , transition functions  $p \in \mathcal{P}$  and rewards  $r \in \mathcal{R}$ . At each timestep  $t$ , the RL agent receives an observation  $o_t$  and generates an action  $a_t$ . The environment then updates the applied action based on the transition function  $p(s_{t+1}|s_t, a_t)$  and assigns a reward to the agent. The goal of the agent is to maximize the cumulative discounted rewards  $R_t =$



Fig. 3. CNN Model: We design a CNN encoder to estimate the state information from image observation. Image augmentation is applied to improve the robustness of state estimation.

$\sum_{t=k}^H \gamma^t r_t$ , where  $H$ ,  $k$ , and  $\gamma$  are the horizon length, current timestep, and the discount factor.

### B. RL by Decision Transformer

While many single-step, forward methods are limited by error accumulation issues, it has been shown that the transformer architecture significantly improves error accumulation [34]. To leverage this, the DT model [28], [29] is adopted in this work as depicted in Fig. 4. To process a sequence of trajectory  $\{\hat{R}_1, s_1, a_1, \dots, \hat{R}_T, s_T, a_T\}$ , the DT model treats the sequence data as *tokens* and embeds them into  $\{z_i\}_{i=1}^{3T}$  with position embedding ( $\hat{R}$  refers to return-to-go [28]). Specifically, the  $i$ -th token, being linearly projected to query  $q_i$ , key  $k_i$ , value  $v_i$ , is then fed into the following self-attention layers. The  $i$ -th output of the self-attention layer is computed by:

$$z_i = \sum_{j=1}^n \text{softmax}(\langle \{q_i, k_{j'}\}_{j'=1}^n \rangle) \cdot v_j$$

Note that the DT model only processes token  $j \in [1, i]$  to predict the following action. At timestep  $t$ , the DT model predicts the action using tokens from the past  $K$  timesteps, where  $K$  is the hyperparameter referred to as the *context length* here. Since the training of DT is supervised, we apply a cross-entropy loss  $\mathcal{L}_{CE}$  to maximize the likelihood of predicted action and ground-truth action from an expert policy. For the online fine-tuning, we add an entropy loss  $\mathcal{L}_{en}$  with empirical weight  $\lambda$  to enhance the exploration ability of our model further ( $\lambda = 0.1$ ) [29]. The final loss is then obtained by:

$$\mathcal{L} = \mathcal{L}_{CE} - \lambda \mathcal{L}_{en}$$

We train our privileged agent from the expert dataset with access to online fine-tuning, as shown in Fig. 2. Similar to [6], we first collect offline datasets by a heuristic expert solution that has access to the full dynamics and states of the cloth to train our privileged DT model. However, the privileged DT model only trained with offline data often has a sub-optimal performance due to the pre-collected dataset distribution. Online fine-tuning is then applied to improve our model performance further [29]. This is implemented by gradually replacing samples from the training buffer with new data collected by the trained agent, as shown in Fig. 2.

### C. Knowledge Distillation for Learning Visual Control Agent

1) *Pre-trained CNN Encoder for State Estimation*: Due to the large domain gap between state and visual observation, learning cloth manipulation directly from visual observation

is particularly challenging in RL. We design a CNN encoder  $F_c$  to encode the observed images for the downstream knowledge distillation framework to minimise the gap between visual and state observation. Specifically, our designed CNN takes an RGB image  $I_t$  as the input and output estimated state representation of the target cloth  $S'_t$ . Since simulation can give access to the state information of cloth, we can easily collect datasets without any manual labelling. In addition, recent work has also shown classic image augmentation techniques can be beneficial to visual RL training [27], [35]. Therefore, we apply image augmentation techniques such as flipping and cropping to improve the performance of the RL agent. Being different from other methods, we only apply image augmentation in the training of the CNN encoder. Instead of augmenting the data in the RL agent training stage, we aim to implicitly improve visual RL agent performance by enhancing the robustness of the pre-trained CNN encoder.

2) *Knowledge Distillation via Weight Initialization*: Without any prior knowledge, training a vision-based RL is rather challenging due to the partial observability of high dynamic cloth states [6]. Therefore, we adopt a KD approach for the learning of vision-based agents. To facilitate this knowledge distillation, we propose an efficient weight initialization-based knowledge distillation framework. This framework aims to train a student policy that can imitate a privilege policy.

The student model employs a similar architecture as the privileged model. The only difference between the student model and the privileged model is that the student policy takes RGB images as input to predict actions. Leveraging pre-trained weight and fine-tuning has shown great potential in the transfer learning of sequence modelling of the Large Language Model (LLM). With the pre-trained weight, the fine-tuned model can obtain not only a boosted performance but also a higher training efficiency [36]. Inspired by that, we initialize the student model by copying the weight of the privileged model to fully leverage the knowledge obtained by the privileged agent. During the distillation procedure, We first apply the pre-trained CNN encoder to embed the high-dimensional image  $I_t$  into low-dimension estimated states  $S'_t$ . The estimated states and reward-to-go are then used to predict the action  $A'_t$ . As shown in Fig.4, with the weight of the encoder and the teacher models frozen, only student policy is updated by minimizing the imitation loss of actions generated by the student and privileged models:

$$s'_t = F_c(I_t), \quad \mathcal{L}_{distill} = \sum_{t=0}^T \mathcal{L}_{CE}(a_t, a'_t)$$

$$(a_0, a_1, \dots, a_T) \sim \mathcal{DT}_{priv}(a_0, s_0, R_0, \dots, a_T, s_T, R_T)$$

$$(a'_0, a'_1, \dots, a'_T) \sim \mathcal{DT}_{stud}(a'_0, s'_0, R_0, \dots, a'_T, s'_T, R'_T)$$

where  $(a_t, s_t, R_t)$  is the state, action and reward-to-go pair in the teacher’s trajectories,  $F_c$  refers to the pre-trained CNN encoder.

## IV. EXPERIMENT

In this section, we present several experiments to validate the effectiveness of our proposed method. We first provide the implementation details and task setup for our approach. We evaluate our method over these tasks to compare results with other SOTA algorithms. Besides, we conduct an ablation study to test the necessity of each component of our method. Finally, we demonstrate our model’s capability to handle noisy state estimation and adapt accordingly. Real-world demonstrations are also conducted in this section.

### A. Tasks and Implementation Details

We evaluate *TraKDis* on three vision-based cloth manipulation tasks. All the simulation experiments are conducted in the SoftGym suite. [3]. For the training of privileged agent training, we use a reduced-state provided by SoftGym. For the student agent, we adopt the observation described by a 32x32x3 RGB environment image demonstrating the end-effector and target cloth. The variation of cloth deformation property is also applied for proper domain randomization of each task.

- *Cloth Fold Image*: The goal of the task is to fold a flattened cloth in half along the edge with two end-effectors. The coordinates of four corners  $(x, y, z)$  are applied as the reduced states. The performance is evaluated by measuring the alignment of two sets of folded corner positions with a penalty based on the displacement of the cloth from its original position.
- *Cloth Fold Diagonal Pinned Image*: The objective of the task is to fold the cloth along a specified diagonal of a square cloth. The coordinates of four corners  $(x, y, z)$  are applied as the reduced states. A heavy block pins one corner of the cloth on the table. The performance is evaluated by measuring the alignment of bottom-left and top-left corner positions.
- *Cloth Fold Diagonal Unpinned Image*: The objective of the task is to fold the cloth along a specified diagonal of a square cloth. The coordinates of four corners  $(x, y, z)$  are applied as the reduced states. The cloth is free to move on the tabletop. The performance is evaluated by measuring the alignment of bottom-left and top-left corner positions.

We report all the results by a normalized version of the performance score provided by SoftGym to compare with other baselines. The normalized performance at timestep  $t$ ,  $s'_t$  is calculated by  $s'_t = (s_t - s_0)/(s_{opt} - s_0)$ , where  $s_{opt}$  is the optimal performance at the task. Following the protocol of [3], [6], we use the normalized performance at the end of the episode.

For the training of CNN, we apply a four-layer convolution neural network (15M) to project the 32x32x3 RGB image into the length of the reduced state. Unlike other RL algorithms, the performance of DT benefits from scaling the model size [29], [30]. Therefore, we adopt a DT model with the embedded dimension of 256, 10 hidden layers and 16 heads, which yields around 40M trainable parameters. After

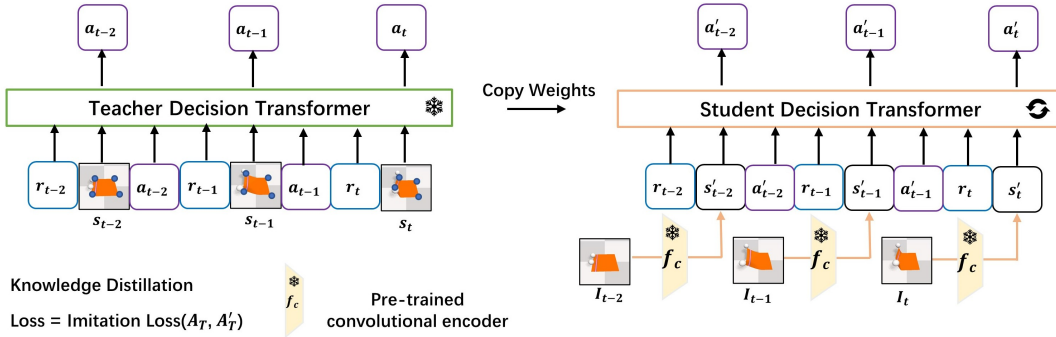


Fig. 4. The figure demonstrates the overview of our knowledge distillation procedure. We run the teacher and student policy simultaneously for knowledge distillation. The teacher policy has access to all the state dynamics of the cloth and thus performs better. The student policy, which receives image inputs, is trained to imitate the actions of the teacher policy. By using a pre-trained CNN, the image can be estimated to encoded input  $s_t$ . Since the student agent and privileged agent have the same architecture, we initialize the weight of the student policy by copying the weights from the teacher’s policy, which aids the knowledge distillation process. The parameters of CNN and teacher policy are frozen during training. Only the student policy is updated via the imitation loss.



Fig. 5. Snapshots of the proposed method. We evaluate our proposed methods by adopting three cloth manipulation tasks. Only image observation is used for the agent to generate trajectory. More demonstration can be found in the multi-media resource.




Fig. 6. The training and validation errors for CNN encoder in three different tasks

generating an action vector, a Tanh function is employed to activate end-effector actions. We optimize our model parameter with the LAMB optimizer [37]. The learning rate is set to  $1e-4$  with a weight decay of  $1e-4$ .

We use the 8000 episodes with a horizon of 100 steps collected from the heuristic algorithm for the offline training of privileged policy and CNN encoder. A total of 200k state-observation-action steps are applied to train privileged policy and CNN for each task. For the student policy, it only takes 10k steps for the image-based policy to converge. We implement our network by Pytorch. All the training is conducted on a machine with Intel Core i9-12900KF and Nvidia GeForce GTX 3080Ti GPU.

## B. Comparison Experiment

We visualize our pre-trained CNN encoder’s training and validation loss in Fig. 6. All the training procedures for three tasks require 20 epochs to be converged. We evaluate our proposed method against several state-of-the-art methods for cloth manipulation to show the advantages of our method in visual robotic cloth manipulation tasks.

- **Teacher:** The teacher DT model trained with privileged cloth states information.
- **DMFD:** image-based state-of-the-art method that combines expert demonstration and RL for deformable object manipulation from [6].
- **Behaviour Cloning:** a behaviour cloning policy that takes image and action pairs for policy training to imitate or replicate behaviours demonstrated by the expert.
- **CURL:** a state-of-the-art image-based RL algorithm that utilizes contrastive representation for RL agent learning.
- **DT-image:** We directly use the RGB image as the input to train the agent. The CNN encoder for state estimation is not applied for training

We use the official implementation from SoftGym and DMFD. All the hyperparameters are kept the same without any modification Table I shows all the results of the comparison experiments. As illustrated in Table I, our methods outperform other baselines in all three image-based tasks. A detailed discussion will be presented in section IV-F.

## C. Ablation Studies

We conduct ablation studies to verify the necessity of each component of our proposed method.

- **Ours:** We pre-train a CNN encoder **with** the image augmentation technique, copy the weight from the teacher agent, and imitate the teacher policy.
- **No-Aug:** We pre-train a CNN encoder **without** image augmentation technique, copy the weight from the teacher agent, and imitate the teacher policy.

TABLE I  
PERFORMANCE OF ROBOTIC CLOTH MANIPULATION TASKS

tasks		Expert-state	Teacher-state	TraKDis(ours)	DMFD	Behavior Cloning	CURL	DT-image
Cloth Fold Image	$\mu \pm \sigma$	0.706	0.740	<b>0.583</b>	0.364	0.159	0.139	0.305
	25 <sup>th</sup> %	$\pm 0.159$	$\pm 0.111$	$\pm 0.190$	$\pm 0.432$	$\pm 0.080$	$\pm 0.312$	$\pm 0.207$
	median	0.637	0.663	<b>0.467</b>	0.000	0.111	0.000	0.162
	75 <sup>th</sup> %	0.726	0.750	<b>0.610</b>	0.517	0.178	0.062	0.299
Cloth Fold Diagonal Pinned Image	$\mu \pm \sigma$	0.907	0.902	<b>0.916</b>	0.778	0.613	0.670	0.469
	25 <sup>th</sup> %	$\pm 0.009$	$\pm 0.004$	$\pm 0.002$	$\pm 0.014$	$\pm 0.008$	$\pm 0.033$	$\pm 0.012$
	median	0.898	0.900	<b>0.916</b>	0.770	0.608	0.670	0.460
	75 <sup>th</sup> %	0.906	0.901	<b>0.916</b>	0.778	0.610	0.677	0.464
Cloth Fold Diagonal Unpinned Image	$\mu \pm \sigma$	0.926	0.921	<b>0.944</b>	0.861	0.663	0.827	0.512
	25 <sup>th</sup> %	$\pm 0.012$	$\pm 0.004$	$\pm 0.004$	$\pm 0.004$	$\pm 0.067$	$\pm 0.034$	$\pm 0.014$
	median	0.917	0.918	<b>0.941</b>	0.859	0.618	0.810	0.505
	75 <sup>th</sup> %	0.923	0.920	<b>0.944</b>	0.862	0.671	0.831	0.510
	75 <sup>th</sup> %	0.937	0.923	<b>0.946</b>	0.863	0.719	0.851	0.521

Comparison Results of Different Model Performances. Expert refers to the oracle with access to the full cloth state and dynamics. The value is obtained from normalized performance defined in the SoftGym environment. The model is obtained from the end of training for each method. We test the model from 100 simulation trials and obtain the mean, standard deviation, 25% and 75% percentiles of performance. For the evaluation, we use a fixed value of -1 and 30 for the initial reward-to-go and context length, respectively. We use the same hyperparameters and pre-trained weights from the official DMFD implementation. DMFD, Behaviour Cloning, and CURL [24] have a model size of 15M, 36M, and 16M, respectively.




Fig. 7. Learning Curves of Ablation Studies. We attain the normalized performance from five random running seeds. We draw the mean performance as the solid lines. Shaded areas stand for the standard deviation of normalized performance. Thanks to the weight initialization, the student model can converge within 60 epochs for all three tasks.

- **No-Pretrain:** We pre-train a CNN encoder with image augmentation technique, **without** copy weight, and imitate the teacher policy.
- **No-Distill:** We remove the distillation part by using the decision transformer and the CNN encoder to imitate the heuristic expert demonstrations.

We draw the learning curves of several variants of our proposed methods as shown in Fig. 7. The results show the effectiveness of every component of our proposed methods. A detailed discussion will be presented in Section IV-F.

#### D. Robustness

To test the robustness of our proposed method, we evaluate our model in a more challenging environment. In this experiment, We introduce a Gaussian noise  $\epsilon$  with a mean of zero and standard deviation  $\sigma = 1e - 3$  to the estimated states  $S'$  for the model training. This aims to test the model's robustness when the CNN has an inaccurate estimation. We use the relatively difficult Cloth Fold task as the case study. The student model is trained for 60 episodes according to the training curves from Fig.7 (a).


The results are demonstrated in Fig. 8. It can be observed that added noise does not degrade the model training and

final performance. We further conduct a comparison experiment of models trained with added noise and without noise. This is achieved by evaluating the model by adding noise that ranges from  $\sigma = 1e - 5$  to  $\sigma = 1e - 3$  in Fig. 8 (b). The results indicate that the model fine-tuned with noise exhibits a higher tolerance for inaccurate state estimation with more significant noise added.


$$S'_{noise} = S' + \epsilon, \quad \epsilon \sim \mathcal{N}(\mu, \sigma^2)$$

#### E. Real-world Experiment

In the real world, we perform the Cloth Fold Unpinned task to test the effectiveness of our proposed method. This is implemented by applying a UR5 robot and a customized parallel gripper, as shown in Fig.9. We rotate the target cloth to different angles to test the performance. One Intel RealSense D435i RGB-D camera is fixed above the scene during the experiment to capture the images. We follow a similar processing procedure in [6] to pre-process our image input to ensure the robustness of our system. No fine-tuning or re-training is performed to fit the real-world setting. Our method is evaluated on a total of ten rollouts.



(a)



(b)

Fig. 8. (a) Training curves with noise added. We train our student model with added Gaussian noise. It can be observed that our model can adapt to the added noise and increase its robustness. (b) Comparison of models. The model trained with added noise has a higher tolerance to added noise. Dashed lines represent the performance where no noise is added. We obtain the result from a mean value of 10 simulation trials.

For the simulation, we obtain a mean result of 0.940. For the real world, we obtain a performance of 0.895. More demonstrations can be found in the multi-media resource.

### F. Discussion

According to the results presented in Table I, our state-based agent demonstrates comparable or superior performance compared to the expert policy. In the case of the image-based environment, *TraKDis* consistently achieves the highest performance across all three tasks. More specifically, our model achieves a performance of 0.583, which corresponds to 78% of the teacher’s performance in the challenging cloth fold task. In addition, our approach achieves performances of 0.916 and 0.944 in the cloth diagonal pinned and unpinned fold tasks, respectively, surpassing both the expert and teacher’s performances. Compared to other baselines, *TraKDis* exhibits an improved performance across three tasks, with enhancements of 21.9%, 13.8%, and 8.3% observed in these respective tasks. Moreover, our algorithm exhibits the lowest standard deviation in terms of model robustness, indicating its reliability and consistency.

By utilizing the image augmentation technique, our pre-trained CNN can employ a more robust state estimation, which stabilizes the training of student policy, as shown in the ablation study. In addition, by initializing the weight from the privileged agent, we significantly improve the training efficiency and performance. The training of our student policy only takes 60 episodes (6000 steps) for the policy to be converged, which is only 0.6% of the total training steps reported in [6]. However, due to the absence of temporal information in the CNN encoder as a state estimator, errors in estimation and inaccuracy regarding action prediction may occur. While this could result in an error accumulation




Fig. 9. Real-world demonstration for cloth diagonal unpinned task

problem for MDP-based methods, we mitigate this issue by using the DT model, adapting the estimated state in a temporal sequence level. Besides, compared to the ablation variants, the weight copy also provides a good initialization for the student policy training. Without the leverage of the teacher’s weight, the model exhibits either a slow performance improvement or performance degradation as shown in the **No-Pretrain** and **No-Distill** ablation variants in Fig.7.

We further evaluate our model in a challenging noisy environment. As shown in Fig.8, It should be observed that our model can adapt to the added noise and converge within 60 epochs. Although adding noise fails to improve the model’s performance, the added noise can be beneficial to the model’s robustness as the standard deviation decreases with more training steps, shown in Fig.8. Results also indicate that by re-training the student model with noise added, our model can obtain a higher tolerance to inaccurate state estimation. Real-world experiments are also conducted to show the effectiveness of our proposed method in the real-world setting.

## V. CONCLUSIONS

In this work, we present a novel approach named *TraKDis* for the learning of visual RL policy for cloth manipulation tasks. While direct training from images is extremely difficult, by leveraging knowledge from the model with access to privileged information, *TraKDis* tries to solve this problem with a one-to-one novel KD approach that leverages weight copy and input state estimation. Although the transformer model has more parameters than MLP-based methods, its sequence modelling capability, considering historical state information for each action, enhances robustness and efficiency in distillation tasks. A comparison experiment is conducted to test the performance respect to state-of-the-art algorithms in three different tasks. The experiment results indicate that *TraKDis* shows better performance in all three tasks. Specifically, our proposed approach demonstrates an increased performance of 21.9%, 13.8%, and 8.3% in these three tasks. Ablation experiments and robustness experiments are also conducted for a more in-depth analysis of our proposed method. To validate our approach further, a real-world experiment is also conducted.

Despite the benefits of the approach, there are some limitations. For instance, our work collects 8000 episodes (80k state-action-reward samples) for privileged policy training. Future work can focus on improving the data efficiency since

the current agent requires a large state and action transition dataset for offline training. Besides, we observe that there is still a 22% performance gap for the distillation of the cloth folding task. Moreover, in contrast to conventional RL approaches, *TraKDis* involves a larger model size. This could potentially lead to issues related to operational speed and memory utilization. Future research could focus on optimizing and reducing the overall size of the model. With recent progress in visual deep RL [38], [39], future work can be devoted towards integrating pre-trained encoder with large-scale image datasets and our distillation approach for more challenging tasks such as complex-shaped cloth (e.g. shirts and trousers) manipulation and multi-goal cloth manipulation tasks.

## REFERENCES

- [1] J. Borràs, “Effective grasping enables successful robot-assisted dressing,” *Science Robotics*, vol. 7, no. 65, p. eabo7229, 2022. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.abo7229>
- [2] J. Matas, S. James, and A. J. Davison, “Sim-to-real reinforcement learning for deformable object manipulation,” in *Conference on Robot Learning*. PMLR, 2018, pp. 734–743.
- [3] X. Lin, Y. Wang, J. Olkin, and D. Held, “Softgym: Benchmarking deep reinforcement learning for deformable object manipulation,” in *Conference on Robot Learning*. PMLR, 2021, pp. 432–448.
- [4] R. Jangir, G. Alenyà, and C. Torras, “Dynamic cloth manipulation with deep reinforcement learning,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 4630–4636.
- [5] Y. Wu, W. Yan, T. Kurutach, L. Pinto, and P. Abbeel, “Learning to manipulate deformable objects without demonstrations,” *arXiv preprint arXiv:1910.13439*, 2019.
- [6] G. Salhotra, I.-C. A. Liu, M. Dominguez-Kuhne, and G. S. Sukhatme, “Learning deformable object manipulation from expert demonstrations,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 8775–8782, 2022.
- [7] S. Levine, C. Finn, T. Darrell, and P. Abbeel, “End-to-end training of deep visuomotor policies,” *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1334–1373, 2016.
- [8] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen, “Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection,” *The International journal of robotics research*, vol. 37, no. 4-5, pp. 421–436, 2018.
- [9] L. Pinto, M. Andrychowicz, P. Welinder, W. Zaremba, and P. Abbeel, “Asymmetric actor critic for image-based robot learning,” *arXiv preprint arXiv:1710.06542*, 2017.
- [10] I.-C. A. Liu, S. Uppal, G. S. Sukhatme, J. J. Lim, P. Englert, and Y. Lee, “Distilling motion planner augmented policies into visual control policies for robot manipulation,” in *Conference on Robot Learning*. PMLR, 2022, pp. 641–650.
- [11] W. Ding, N. Majcherczyk, M. Deshpande, X. Qi, D. Zhao, R. Madhivanan, and A. Sen, “Learning to view: Decision transformers for active object detection,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 7140–7146.
- [12] S. K. Ramakrishnan, D. Jayaraman, and K. Grauman, “An exploration of embodied visual exploration,” *International Journal of Computer Vision*, vol. 129, pp. 1616–1649, 2021.
- [13] W. Ding, N. Majcherczyk, M. Deshpande, X. Qi, D. Zhao, R. Madhivanan, and A. Sen, “Learning to view: Decision transformers for active object detection,” *arXiv preprint arXiv:2301.09544*, 2023.
- [14] S. Humeau, K. Shuster, M.-A. Lachaux, and J. Weston, “Polyencoders: Transformer architectures and pre-training strategies for fast and accurate multi-sentence scoring,” *arXiv preprint arXiv:1905.01969*, 2019.
- [15] Z. Zhu, K. Lin, A. K. Jain, and J. Zhou, “Transfer learning in deep reinforcement learning: A survey,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [16] J. Borràs, G. Alenyà, and C. Torras, “A grasping-centered analysis for cloth manipulation,” *IEEE Transactions on Robotics*, vol. 36, no. 3, pp. 924–936, 2020.
- [17] I. Garcia-Camacho, M. Lippi, M. C. Welle, H. Yin, R. Antonova, A. Varava, J. Borras, C. Torras, A. Marino, G. Alenyà, and D. Kragic, “Benchmarking bimanual cloth manipulation,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1111–1118, 2020.
- [18] D. Seita, N. Jamali, M. Laskey, A. K. Tanwani, R. Berenstein, P. Baskaran, S. Iba, J. Canny, and K. Goldberg, “Deep Transfer Learning of Pick Points on Fabric for Robot Bed-Making,” in *International Symposium on Robotics Research (ISRR)*, 2019.
- [19] J. Qian, T. Weng, L. Zhang, B. Okorn, and D. Held, “Cloth region segmentation for robust grasp selection,” in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 9553–9560.
- [20] H. Ha and S. Song, “Flingbot: The unreasonable effectiveness of dynamic manipulation for cloth unfolding,” in *Conference on Robotic Learning (CoRL)*, 2021.
- [21] Z. Xu, C. Chi, B. Burchfiel, E. Cousineau, S. Feng, and S. Song, “Dexterity: Deformable manipulation can be a breeze,” in *Proceedings of Robotics: Science and Systems (RSS)*, 2022.
- [22] W. Chen, D. Lee, D. Chappell, and N. Rojas, “Learning to grasp clothing structural regions for garment manipulation tasks,” *arXiv preprint arXiv:2306.14553*, 2023.
- [23] Y. Wang, Z. Sun, Z. Erickson, and D. Held, “One policy to dress them all: Learning to dress people with diverse poses and garments,” in *Robotics: Science and Systems (RSS)*, 2023.
- [24] M. Laskin, A. Srinivas, and P. Abbeel, “Curl: Contrastive unsupervised representations for reinforcement learning,” in *International Conference on Machine Learning*. PMLR, 2020, pp. 5639–5650.
- [25] D. Hafner, T. Lillicrap, I. Fischer, R. Villegas, D. Ha, H. Lee, and J. Davidson, “Learning latent dynamics for planning from pixels,” in *International conference on machine learning*. PMLR, 2019, pp. 2555–2565.
- [26] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
- [27] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, et al., “An image is worth 16x16 words: Transformers for image recognition at scale,” *arXiv preprint arXiv:2010.11929*, 2020.
- [28] L. Chen, K. Lu, A. Rajeswaran, K. Lee, A. Grover, M. Laskin, P. Abbeel, A. Srinivas, and I. Mordatch, “Decision transformer: Reinforcement learning via sequence modeling,” *Advances in neural information processing systems*, vol. 34, pp. 15084–15097, 2021.
- [29] Q. Zheng, A. Zhang, and A. Grover, “Online decision transformer,” in *International Conference on Machine Learning*. PMLR, 2022, pp. 27042–27059.
- [30] K.-H. Lee, O. Nachum, M. S. Yang, L. Lee, D. Freeman, S. Guadarrama, I. Fischer, W. Xu, E. Jang, H. Michalewski, et al., “Multi-game decision transformers,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 27921–27936, 2022.
- [31] G. Hinton, O. Vinyals, and J. Dean, “Distilling the knowledge in a neural network,” *stat*, vol. 1050, p. 9, 2015.
- [32] A. Kumar, Z. Fu, D. Pathak, and J. Malik, “Rma: Rapid motor adaptation for legged robots,” 2021.
- [33] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, “Learning quadrupedal locomotion over challenging terrain,” *Science robotics*, vol. 5, no. 47, p. eabc5986, 2020.
- [34] M. Janner, Q. Li, and S. Levine, “Offline reinforcement learning as one big sequence modeling problem,” *Advances in neural information processing systems*, vol. 34, pp. 1273–1286, 2021.
- [35] M. Laskin, K. Lee, A. Stooke, L. Pinto, P. Abbeel, and A. Srinivas, “Reinforcement learning with augmented data,” *Advances in neural information processing systems*, vol. 33, pp. 19884–19895, 2020.
- [36] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [37] Y. You, J. Li, S. Reddi, J. Hseu, S. Kumar, S. Bhojanapalli, X. Song, J. Demmel, K. Keutzer, and C.-J. Hsieh, “Large batch optimization for deep learning: Training bert in 76 minutes,” *arXiv preprint arXiv:1904.00962*, 2019.
- [38] R. M. Shah and V. Kumar, “Rrl: Resnet as representation for reinforcement learning,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 9465–9476.
- [39] C. Wang, X. Luo, K. Ross, and D. Li, “Vrl3: A data-driven framework for visual deep reinforcement learning,” *Advances in Neural Information Processing Systems*, vol. 35, pp. 32974–32988, 2022.