

ClothPPO: A Proximal Policy Optimization Enhancing Framework for Robotic Cloth Manipulation with Observation-Aligned Action Spaces

Libing Yang¹, Yang Li^{1*} and Long Chen²

¹East China Normal University

²The Hong Kong University of Science and Technology

51255901139@stu.ecnu.edu.cn, yli@cs.ecnu.edu.cn, longchen@ust.hk

Abstract

Vision-based robotic cloth unfolding has made great progress recently. However, prior works predominantly rely on value learning and have not fully explored policy-based techniques. Recently, the success of reinforcement learning on the large language model has shown that the policy gradient algorithm can enhance policy with huge action space. In this paper, we introduce ClothPPO, a framework that employs a policy gradient algorithm based on actor-critic architecture to enhance a pre-trained model with huge 10^6 action spaces aligned with observation in the task of unfolding clothes. To this end, we redefine the cloth manipulation problem as a partially observable Markov decision process. A supervised pre-training stage is employed to train a baseline model of our policy. In the second stage, the Proximal Policy Optimization (PPO) is utilized to guide the supervised model within the observation-aligned action space. By optimizing and updating the strategy, our proposed method increases the garment's surface area for cloth unfolding under the soft-body manipulation task. Experimental results show that our proposed framework can further improve the unfolding performance of other state-of-the-art methods. Our project is available at <https://vpx-ecnu.github.io/ClothPPO-website/>.

1 Introduction

With the continuous technological advancements in the field of embodied intelligence [Gupta *et al.*, 2021], emerging application scenarios for robots are gradually coming to the fore, such as service robots. The manipulation of flexible objects is a fundamental and common skill in human domestic labor, including activities such as unfolding and tidying. A robot capable of manipulating clothes could alleviate the burden of household chores and significantly enhance the practicality of robots in our lives.

Compared to rigid body manipulation [Zeng *et al.*, 2020], predicting morphological changes in fabrics through visual

inputs is very difficult due to their high dimensionality and complex dynamics [Seita *et al.*, 2021; Lin *et al.*, 2022] in cloth manipulation. The reliance on supervised learning for such tasks [Xiong *et al.*, 2023] demands an extensive collection of labeled training data, which can be a time-intensive and arduous process. [Matas *et al.*, 2018] proposes to use self-supervised learning to deal with the manual labeling problem. However, supervised learning models often exhibit limited generalization capabilities [Nair *et al.*, 2017]. They tend to excel in conditions similar to their training environments but may falter in new or unforeseen scenarios. This characteristic becomes particularly problematic in robotic manipulation.

The emergence of Deep Reinforcement Learning (DRL) [François-Lavet *et al.*, 2018] offers a new paradigm to achieve more generality and adaptability. Previous methods [Ha and Song, 2021; Canberk *et al.*, 2022; Avigal *et al.*, 2022] employ a value function to predict the unfolded coverage in the task of folding clothes and achieve very impressive results. By employing supervised learning on the coverage area values directly in the observation [Wu *et al.*, 2020a], these methods successfully build the robot's ability to increase the performance of unfolding garments. However, these methods inherit self-supervision from the DRL framework but do not use multi-step reward, still highly rely on the accuracy of the value function. This makes these methods inherit the problems of supervised learning.

On the other hand, policy-based reinforcement methods [Chen* *et al.*, 2023] have shown mediocre performance in practice. Traditional action commands in policy-based DRL are typically represented by low-dimensional discrete or continuous vectors. For instance, robot actions are often controlled by specifying directions (forward, backward, left, right) and movement distances. This results in an inherently discrete low-dimensional action space, compared with the high-dimensional nature of visual input space. [Wu *et al.*, 2020b] employ CNN-based networks and MLPs to map the visual inputs to this low-dimensional action space. As the essential complexity of the garment's dynamics, this mapping potentially leads to significant information loss, which hinders the learning of policy. However, the dual-pose action directly within the pixel space meets another problem an exponential escalation in the number of feasible actions. The pixels in the image as a potential point of action. Considering an image dimension of 256, the resulting potential action

*corresponding author

choices would exponentially rise to the magnitude of 10^{10} , which is impossible to solve.

In this paper, we propose ClothPPO to fully take advantage of reinforcement learning for cloth unfolding tasks. To this end, an observation alignment policy is introduced to map the image directly to a high-dimension action distribution similar to the spatial action maps [Wu *et al.*, 2020a]. To address the information loss from images to low-dimensional actions and directly connect image pixels with robot actions, we first to use spatial action maps in policy-based RL to model policy. This makes our proposed policy have enough capability to capture complex dynamics of fabric changes during manipulation. In addition, inspired by the success of reinforcement learning from human feedback [Ouyang *et al.*, 2022] on large language model, a two-stage learning framework with PPO [Schulman *et al.*, 2017] is proposed to solve the large action space learning problem [Ramamurthy *et al.*, 2022] in policy-based RL. At first, we offline pre-train a UNet-based policy in a self-supervised learning fashion. Then an actor-critic architecture [Schulman *et al.*, 2017] is employed with PPO to further online fine-tune the proposed policy.

We have the following three contributions. 1) Our method generates actions by sampling from a distribution-modeled policy, encouraging thorough exploration of the action space without relying on accurate estimations. Our work is the first to use a policy-based RL approach to directly select robot actions in the huge pixel space aligned with observation. 2) The vast action space makes the exploration process of the RL inefficient and difficult to learn. Inspired by the training process of ChatGPT, we bridge value-based and policy-based approaches, using the value-based method for initialization and employing PPO to enhance the pre-trained model’s effectiveness in addressing the issue of large action spaces in RL. 3) We apply this method to the complex task of robotic cloth unfolding which achieves SOTA results on this difficult practical application problem, and contribute an environment based on our method to the robotics community.

2 Related Work

Research in the domain of flexible cloth manipulation has extensively focused on vision-based methodologies to orchestrate manipulative actions for an array of tasks, encompassing folding and unfolding. In the nascent stages of robotic garment folding, prevalent practices entailed initially unfurling the cloth into predefined, antecedently known states, subsequently engaging in the automated folding via procedural or rudimentary rule-based heuristic techniques [Triantafyllou *et al.*, 2016]. The perception module mainly uses CNN [Avigal *et al.*, 2022; Hu *et al.*, 2023], Transformer [Mo *et al.*, 2022; Chen *et al.*, 2024] and other network structures.

In robotics cloth manipulation, the primitives used for manipulation are categorized into two types: quasi-static and dynamic. Previous work has actively explored the determination of parameters for predefined operational primitives [Mulero *et al.*, 2023]. Dynamic methods calculate high-speed dynamic throwing actions directly from overhead images, which can smooth a garment to 80% coverage within an average of three actions [Ha and Song, 2021].

Current mainstream approaches are divided into model-based learning, which models the environment’s dynamics while neglecting policy learning, and model-free learning, which derives cloth manipulation strategies from images through supervised or reinforcement learning. The latter faces challenges in achieving high robustness due to ambiguities in 2D observational data, insufficiently flexible kinematic structures, and complex environmental interactions. Some papers address this issue by learning from demonstrations [Xiong *et al.*, 2023]. Demonstrations can be obtained from scripted action sequences based on heuristic experts [Doumanoglou *et al.*, 2016] or cloth descriptors [Ganapathi *et al.*, 2021]. [Avigal *et al.*, 2022] engage in self-supervised learning augmented by minimal expert demonstrations to execute cloth smoothing and folding with an assortment of manipulation primitives, achieving a folding success rate as high as 93% within 120 seconds from a randomized initial configuration. However, this method is trained directly in the real world, has weak generalization, requires costly manual demonstration annotations, and has high training costs. Another learning approach is model-free reinforcement learning, [Canberk *et al.*, 2022] introduces a self-supervised reward-based approach for learning cloth normalization and alignment, diminishing the complexity of specific task-oriented downstream policies. However, due to the cloth’s potential for various wrinkle structures, it is challenging to generalize to different target situations. There is also work primarily based on models, suggesting methods for modeling the dynamics of flexible objects for manipulation. [Lin *et al.*, 2022] employing a particle-based GNN dynamic model informed by point cloud observations. Another study [Huang *et al.*, 2023] proposes estimating the real-world cloth 3D mesh under action conditions. Yet, learning accurate 3D representations of real flexible objects within an end-to-end framework remains difficult, the effectiveness of such methods is heavily reliant on the performance of 3D perceptual dynamics modeling.

3 Problem Formulation

3.1 Markov Decision Process

The task of cloth unfolding presents a complex challenge, where robots must interact with fabric over discrete time steps t to maximize coverage area. This necessitates devising an effective manipulation strategy to explore the problem of learning robotic actions a_t from visual observations o_t . A key difficulty lies in the inability to fully observe the cloth’s state in real-world scenarios. It relies on partial visual inputs, creating a gap between perception and complete state awareness. Therefore, we model cloth manipulation as a Partially Observable Markov Decision Process (POMDP) [Sondik, 1971]. A POMDP is defined by the tuple $(\mathcal{S}, \mathcal{A}, \mathcal{O}, P, O, R, \gamma)$, where \mathcal{S} denotes the state space of all possible configurations of the cloth. \mathcal{A} represents the action space, comprising all the manipulative actions a robot can perform on the cloth. \mathcal{O} is the observation space, reflecting the visual observations o_t available to the robot at time t . $P(s'|s, a)$ is the state transition probability function that provides the probability of moving to state s' from state s af-

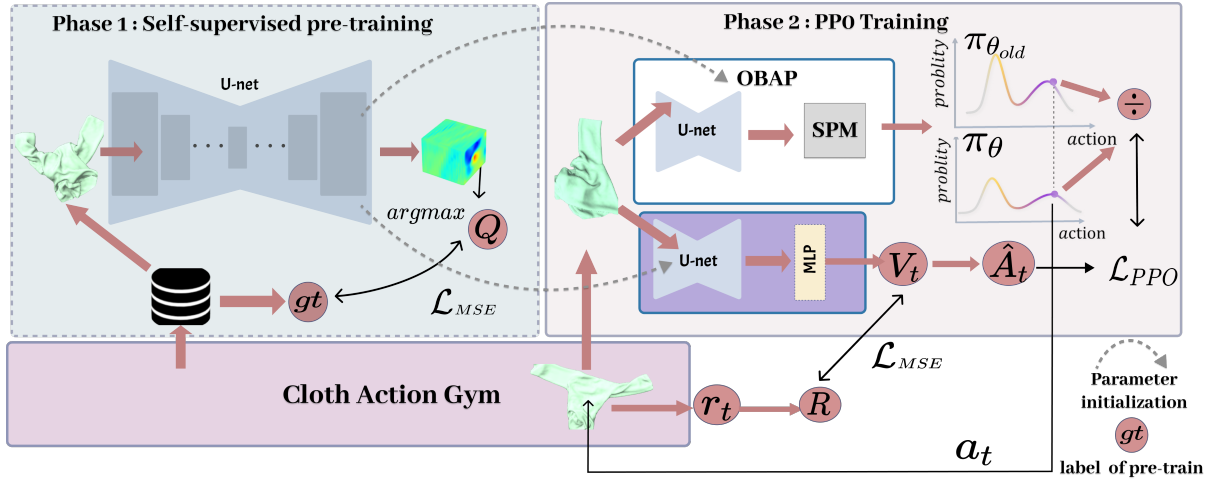


Figure 1: **Overview of ClothPPO.** The first phase [Canberk *et al.*, 2022] involves self-supervised pre-training, which uses data from repeated actions in the environment to collect labels. The model estimates canonicalized alignment grasping label and selects the maximum estimated value to output the action. We introduce a long-term reward mechanism to improve the model’s performance in the PPO training phase.

ter taking action a_t . $O(o|s', a)$ is the observation function that gives the probability of receiving an observation o after taking an action a_t and ending up in state s' . $R(s, a)$ is the reward function that assigns a numeric reward for taking an action a_t in state s , guiding the robot to maximize the cloth’s coverage area. $\gamma \in [0, 1]$ is the discount factor that values the importance of immediate versus future rewards. Suppose the strategy parameter is θ and our optimization objective is to maximize the average return. A reward is $r_t = r(o_t, a_t, o_{t+1})$. Given the sequence tuples of (o_t, a_t, r_t) . The optimization objective can be expressed as $J(\theta) = E[R]$, where $J(\theta)$ is the objective function of the strategy and $R = \sum_{t=0}^T \gamma^t r_t$:

$$\pi_\theta = \operatorname{argmax}_\pi \mathbb{E}_\pi \left[\sum_{t=0}^T \gamma^t R(o_t, a_t) \right], \quad (1)$$

It is to find a policy $\pi_\theta(a_t|o_t)$ that maximizes the expected sum of discounted future rewards. This involves calculating an expected value that balances immediate and future rewards. To increase the probabilities of taking actions in high-reward trajectories. In this way, during the training process, the policy gradually adjusts its parameters to improve its performance.

3.2 Definition of Observation and Action

In the task of clothing unfolding, the agent is assigned to grasp the points on the cloth and unfold it into a flat state. This task can be accomplished through a series of dual-pose motion primitives, such as pick and place. Given the original header observation, policy decides the optimal actions $a_t = (a_{p1}, a_{p2}) \in \mathcal{A}$. Observation is an RGBD image $o_t \in \mathbb{R}^{W \times H \times C}$, where H and W are the height and width dimensions of the image. The function that maps observations to actions is defined as:

$$\pi_\theta(a_{p1}, a_{p2}|o_t) = p(a_{p1}|o_t) \cdot p(a_{p2}|o_t, a_{p1}), \quad (2)$$

where $p(a_{p1}|o_t)$ and $p(a_{p2}|o_t, a_{p1})$ are the probabilities of these actions given the observation o_t . Since observations can

be projected into a 2D-pixel space, with the third dimension held at a fixed depth, it treats this as a two-dimensional action. This constrains a_{p1} and a_{p2} to two points in pixel coordinate.

The dual-pose joint action model, which incorporates both direct and visual input modalities, is formalized as a multivariate discrete scalar action space. This space comprises four variables amalgamated within the pixel space. However, the two-step operational action space involved in making decisions directly at the image pixel level is huge.

4 Method

As shown in Fig.1, the training process of our proposed method consists of two phases, our contributions focus on the second phase:

Self Supervised Pre-training. The model is founded on self-supervised principles and adopts a regression methodology for individual pixel values Q using Mean Squared Error (MSE) loss. However, its inherently step-by-step prediction approach presents a key limitation: a noticeable absence of long-term guidance.

PPO Training. In the second phase, we use the PPO to fine-tune the pre-trained model. In this phase, we introduce a long-term reward mechanism to improve the model’s performance in goal-oriented tasks. The PPO iteratively optimizes the model’s policy to better adapt to the goal task and guides the model’s learning process through reward signals. With these two phases of training, our goal is to enable the model to generate excellent grasping actions in grasping tasks and show higher performance in goal-oriented tasks. This two-stage training strategy improves the generalization ability and task adaptability of the model.

4.1 Observation Alignment Policy

Redefine Action to Reduce the Action Space. As formulated Eq. 2, the dual-pose joint action comprises four variables amalgamated within the pixel space. We redefine the

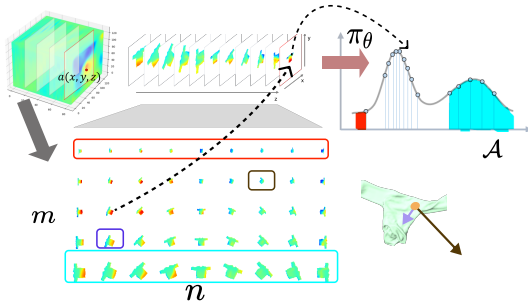


Figure 2: SPM: Our action spaces and sampling action using spatial policy maps. The series of smaller maps at the bottom are different slices of the spatial policy maps, each representing a layer with a mask at a different scale and rotation. The masks applied to each layer filter out invalid actions that would result in the robot’s end-effector interacting with areas beyond the cloth. The variation in the sizes of the masks corresponds to different scales, affecting the granularity of actions that can be sampled.

action space of the above PDMDP as another multivariate discrete action space and re-describe the joint action using the starting point and the line segment originating from it. This effectively reduces the dimensionality of the action from a larger two-step pixel space action to a one-step pixel 10^6 action space, while ensuring the correlation between the two poses. The action at time t can be presented as

$$a_t = (x, y, \phi, d) = (x, y, z) \in \mathcal{A}, \quad (3)$$

where each element represents a different dimension of the action space, x, y represents the set of action positions in pixel coordinate, n is the number of a set of angles $\Phi = \{\phi_i | \phi_i = \frac{360^\circ * i}{n}, i = 1, 2, \dots, n\}$. m represents the number of a set of movement distances $D = \{d_i | d_i = \frac{b-a}{m} * i + a, i = 1, 2, \dots, m\}$, where b, a represent the max and min movement distance. $z \in \Phi \times D$ denotes the combination of an angle and a movement distance.

The line segments on the clothes explain this definition in Fig.2. The black and purple line segments represent two different actions.

Rotate and Scale Observation. For intuitive understanding, we refer to the setting of spatial action maps [Wu *et al.*, 2020a] to rotate and scale the observation to align action space with the visual input, x, y, z is the position of the 3D map. The corresponding joint action is determined by the index of the pixel in the 3D map, so a single pixel can correspond to an action. Spatial action maps facilitate the network’s ability to learn and predict the Q value for each point, guided by the characteristics of the state representation. They are used to predict the value function $Q(s_t, a_t)$, with the map value representing the expected reward for choosing a particular action. We use it as spatial policy maps to model the policy itself, which represents the preference for choosing various possible actions in a given state. The value in our map is the probability of opting for a particular action. As far as we know, ours is the first work to use it in policy-based RL.

4.2 Sampling Action from OBAP

Flatten 3D Policy Maps to Categorical Distribution. Spatial policy maps output from the network indicate the probability of each potential action within the environment. To sample an action a_t according to the policy $\pi_\theta(\cdot | o_t)$, the multi-dimensional spatial policy maps must be converted into a one-dimensional logits vector, which represents the log probabilities before normalization. This transformation process is referred to as flattening.

When employing a categorical action spaces [Tutz, 1991], the output from the network can be transformed into a discrete probability distribution through a softmax operation on the logits vector, denoted as \mathbf{z}_t :

$$p(a_t | o_t) = \frac{e^{\mathbf{z}_t, a_t}}{\sum_k e^{\mathbf{z}_t, k}}, \quad (4)$$

where a_t is the action taken at time t and k iterates over all possible actions.

Spatial Policy Mask. Sampling actions directly on the existing spatial policy maps can lead to many invalid actions (e.g., futile grasps or actions exceeding designated limits). Since the map denotes the probability of selecting an action, the corresponding mask represents regions with non-zero selection probabilities, visually equivalent to clothing parts in the observation. Due to the different rotations and scalings of each layer, the shape of the cloth mask is also different (Fig. 2). For the mask, these clothing areas are assigned true values while other areas are marked as false values, which are displayed in black in the simulation. We have added a different mask to each layer of the spatial policy maps. For each state observation o_t at time t , the applicable action space is obtained by excluding the invalid actions identified by the mask M_t . As indicated by the equation below, we have incorporated a mask into the action space on top of PPO to filter out futile actions:

$$\pi_\theta^*(a_t | o_t) = \begin{cases} \pi_\theta(a_t | o_t) & \text{if } M_t(a_t) = 1, \\ -\text{inf} & \text{if } M_t(a_t) = 0, \end{cases} \quad (5)$$

where $M_t(a_t)$ is the masking function at time t that returns 1 for valid actions and 0 for invalid actions.

As shown in the red area in Fig. 2, if the observation is scaled to smaller, from an action perspective, the line segment will be relatively longer, and the movement will be further. From the perspective of observation, there are fewer pixels left after being masked. Then its proportion in the distribution will become smaller after flattening. In other words, violent action will be a small probability to sample, and vice versa (blue area). The benefit of this Spatial Policy Mask(SPM) is the probability of actions of different granularities being sampled can be controlled by the size of the mask corresponding to the scaling, thus enabling a more efficient search through the action space for the optimal set of movements.

In this way, we get our Observation Alignment Policy(OBAP). To account for constraints in the action space, mask vector \mathbf{m}_t is applied element-wise to the logits before the softmax operation. This mask has binary entries where 1 indicates a valid action and 0 is an invalid one. The probability of selecting any specific action a_t from the masked action

space is given by $\tilde{p}(a_t|o_t)$, which respects the feasibility of actions as determined by the constraints encoded in the mask \mathbf{m}_t .

$$\tilde{p}(a_t|o_t) = \frac{e^{\mathbf{z}_t \cdot \mathbf{m}_t \cdot a_t}}{\sum_k e^{\mathbf{z}_t \cdot \mathbf{m}_t \cdot k}}, \quad (6)$$

where a_t is the action taken at time t and k iterates over all possible actions, \mathbf{z}_t is the logits vector which is transformed into a discrete probability distribution through a softmax operation.

4.3 Proximal Policy Optimization Finetuning

By employing the PPO [Schulman *et al.*, 2017], we have augmented the efficacy of the supervised policy in cloth manipulation tasks.

Throughout the training, an action from the action space is sampled at each timestep. Specifically, we sample an action a_t according to the policy $\pi_\theta(a_t|o_t)$ derived from the current model parameters θ . The core of the PPO update is the clipped surrogate objective, which is used to modify the parameters θ of the policy network to maximize the expected return. Given a set of trajectories (rollouts), collected under the current policy $\pi_{\theta_{old}}$, PPO updates the policy parameters by optimizing the following loss function:

$$L^{PPO}(\theta) = \mathbb{E}[\min(p_t(\theta)\hat{A}_t, \text{clip}(p_t(\theta), 1 - \epsilon, 1 + \epsilon)\hat{A}_t)], \quad (7)$$

where $\hat{A}_t = r_t + \gamma V_t(o_{t+1}) - V_t(o_t)$, is an estimator of the advantage function at timestep t , $p_t(\theta)$ denotes the probability ratio $p_t(\theta) = \frac{\pi_\theta(a_t|o_t)}{\pi_{\theta_{old}}(a_t|o_t)}$, and ϵ is a hyperparameter, typically a small constant like 0.1 or 0.2, that defines the clipping range to avoid overly large policy updates. The advantage function estimator \hat{A}_t quantifies the benefit of taking action a_t in o_t over the expected value of the policy at that time. It is calculated using a value function $V(o_t)$. $\mathbb{E}[\cdot]$ is taken over the distribution of actions and states encountered under the current policy $\pi_{\theta_{old}}$.

Model Architecture. As shown in Fig. 1, the PPO phase fine-tuned the OBAP in simulation via an actor-critic architecture [Schulman *et al.*, 2017].

Actor. OBAP selects actions based on the current observation. The goal of the actor is to maximize the cumulative reward. OBAP accepts RGBD observations o_t as input and outputs per-step action a_t . OBAP gets the next observation o_t by interacting with the environment with a_t and optimizes its parameters based on the multi-step reward to enhance the policy’s performance. In the initial phase of PPO training, OBAP is a copy of the pre-train model. As the PPO finetunes proceed, the parameters of OBAP are iteratively updated through the policy gradient (see Eq. 7).

Critic. The critic is the value function $V(o_t)$ (see Eq. 7) estimator that evaluates the performance of OBAP. The label of the critic is calculated by multi-step rewards $R(o_t, a_t)$ (see Eq. 1). The critic accepts o_t as input and outputs an estimate of the corresponding $R(o_t, a_t)$. Critic obtains reward signals by interacting with the environment and uses these signals to update the parameters of the value function. The network consists of a pre-trained U-Net combined one-layer MLP, that is updated using MSE loss for optimization.

4.4 Training Details

Reward Design. In addressing the cloth unfolding task, our objective is to maximize the cloth’s coverage over the workspace, referred to as Coverage. To effectively guide the agent toward achieving higher coverage, we formulate a reward function that integrates both sparse and dense rewards. The reward at each time step r_t , is determined by assessing changes in coverage. This change is calculated as the ratio of the difference in current and previous coverage to the maximum possible coverage. The agent will receive a sparse reward of $r_t = -1$ if the coverage decreases, and a dense reward of $r_t = \Delta_{Coverage} \times 20$ if the coverage increases. Additionally, achieving a Coverage greater than 0.9 triggers a sparse reward of 5, further motivating the agent towards optimal cloth manipulation.

Reward Scaling. Given the presence of abrupt, sparse rewards in the reward distribution, overly large numbers can cause instability during network training. Therefore, we have scaled the rewards to act as normalization, stabilizing the training. The scale is realized by dividing the reward r_t by the standard deviation of the reward sequence RS . This is mathematically reflected in the following equation:

$$\tilde{r}_t = \frac{r_t}{\sigma(RS)}, \quad (8)$$

where \tilde{r}_t represents the normalized reward at time t , and $\sigma(RS)$ denotes the standard deviation of the reward sequence RS . By applying this operation, the rewards retain their relative differences in value but are expressed in measures of standard deviations from the mean of the reward distribution.

5 Experiments Settings

Environment and Tasks. Referring to the setting of [Camberk *et al.*, 2022], the simulation environment is built on top of PyFleX [Li *et al.*, 2018] bindings with Nvidia FleX bindings provided by SoftGym [Lin *et al.*, 2020]. We randomly sample the initial state of the cloth from a subset of Cloth3D [Bertiche *et al.*, 2020], which contains filtered meshes of various garments. We only trained on the T-shirt task and tested on five different tasks of 200 wrinkled clothing with low coverage. To evaluate our policy, we load tasks from the testing task datasets and then run the policy for 10 steps or until the cloth is out of observation or 10 steps end.

Metrics. In the simulation, the workspace is discretized into a grid. Each cell on the grid is either covered by cloth or not. We define a coverage metric C_{sim} , which is the sum of the areas of all grid cells covered by the cloth: $C_{sim} = \sum_{i=1}^N A_i$, where N is the number of cells covered by the cloth and A_i is the area of the i -th cell. To evaluate the cloth manipulation performance, we compute the coverage percentage C_{pct} , which is the ratio of C_{sim} to the total area A_{flat} designated for the task, multiplied by 100: $C_{pct} = \frac{C_{sim}}{A_{flat}} \times 100\%$. Final coverage measures the coverage at the end of an episode and delta coverage is final coverage minus initial coverage. Percent positive is the proportion of actions that increased coverage.

	Shirt	Pants	Jumpsuit	Skirt	Dress
Flingbot	53.64 / 00.15 / 51.81	65.66 / 01.24 / 51.65	53.15 / 00.07 / 49.62	56.26 / 00.78 / 52.32	66.38 / 01.07 / 51.27
PPO	50.48 / 07.34 / 52.96	52.28 / 01.97 / 52.14	50.24 / -0.66 / 51.61	49.07 / 14.45 / 53.00	52.54 / -18.91 / 53.78
Cloth Funnels	86.91 / 40.98 / 71.10	70.81 / 20.60 / 54.22	73.13 / 21.45 / 62.89	77.99 / 40.82 / 56.09	79.15 / 33.23 / 55.40
Cloth Funnels-P	73.08 / 29.57 / 72.14	71.77 / 22.77 / 55.62	73.76 / 23.03 / 72.68	69.70 / 35.44 / 66.14	76.44 / 30.43 / 68.52
ClothPPO	80.24 / 34.34 / 73.84	80.58 / 30.38 / 72.14	83.04 / 28.10 / 77.17	85.74 / 45.17 / 74.22	86.56 / 40.39 / 73.36


 Table 1: Comparison with baseline methods across various clothing categories Final Coverage \uparrow /Delta Coverage \uparrow /Percent Positive \uparrow .


Figure 3: Visualizations of same episodes of before and after PPO finetuning in the same task. The cloth starts in an unspread, crumpled state and progresses through different stages towards being more unfolded and spread out. Across the sequence, one can observe changes in the cloth’s position and shape, indicating the unfolding action taken by the agent.

Cloth Action Gym. We build a reinforcement learning environment Cloth Action Gym that is easy to use and extend by wrapping an existing simulation with the OpenAI Gym [Brockman *et al.*, 2016] API. We also apply Tianshou [Weng *et al.*, 2022]’s PPO algorithm to implement the above method. In this way, others can quickly modify the existing task environments and high-level action primitives to suit their research goals, and use any reinforcement learning frameworks that are compatible with our environment.

Baselines. We compare our method with the following baselines on various categories of clothing: Flingbot [Ha and Song, 2021], which utilizes the dynamic dual-arm Fling action primitive, leveraging coverage as a label in its value map. We employ the original model and action primitive configurations provided by Flingbot, which is trained on rectangular fabric of varying sizes. Experiments are conducted on multi-category clothing tasks. PPO means the PPO algorithm without pre-training model. Cloth Funnels [Canberk *et al.*, 2022], a multi-action primitive method for dual-arm robots using canonicalized-alignment labels that requires shape priors. Cloth Funnels-P, Cloth Funnels, which solely employs the Pick-Place action in line with our action, is also the pre-trained model of ours ClothPPO.

6 Results

Comparing with Baselines. Table.1 shows the performance comparison of various methods applied to five differ-

ent types of clothing tasks. Each entry in the table shows three numbers separated by a slash representing three different performance metrics. Apart from the T-shirt Task, ClothPPO has attained superior results across all other metrics. This is because the shape of prior labels related to clothing categories used by Cloth Funnels is particularly prominent in this Task. But our ClothPPO still performs better on other types of cloth tasks than the method that requires the use of additional robotic arms to operate Fling action primitives. The most compelling evidence of its efficacy is the significant improvement it has shown when compared to the pre-training model. This clearly illustrates the effectiveness of our approach. This makes perfect sense considering our rewards are directly tied to enhance these three metrics.

Generalize to Unseen Cloth Types. The reward system in ClothFunnels relies on the alignment of cloth shapes, thereby requiring knowledge of the fabric’s geometry beforehand. Consequently, optimal outcomes for different garment shapes necessitate distinct training models. However, in practical terms, this approach is highly inefficient. Conversely, ClothPPO’s requisite reward is based on the coverage of the cloth, which is a universal metric applicable to all clothing categories, thereby obviating the need for training individual models for each distinct class of garments. Our experiments demonstrate that ClothPPO, even when exclusively trained on the Shirt task, exhibits excellent performance and adaptability across multiple types of clothing. Notably, ClothPPO’s




Figure 4: Reward Ablation. We compare three rewards: Threshold Achievement Reward (pink line): ends an episode when coverage exceeds 0.95. Its design enhances computational efficiency and motivates the model to complete tasks quickly. Over Achievement Reward (blue line): Despite achieving a coverage greater than 0.95, this function continues the episode. Immediate Termination Reward (green line): provides a reward and ends the episode as soon as coverage surpasses 0.95. Shaded areas show differences in multiple training experiments

ability to retain high-performance metrics across varying garment types indicates a substantial capacity for generalization. This is crucial for practical applications wherein robots may be tasked to handle a wide array of garments.

Comparing Performance Before and After PPO Training.

Fig. 3 demonstrates a comparison of performance on a consistent task set by both the pre-training model and ClothPPO. The presence of robotic grippers highlighted by red lines, suggests an interaction between the robotic manipulator and the cloth. Within the same task, ClothPPO displays a more significant improvement in coverage at every step. In the final stage of the pre-train model, the coverage stagnates due to the action being void. The pre-training model approaches reliance on the *argmax* action selection strategy, thereby limiting its exploration of other potential actions. In contrast, ClothPPO manifests a broader scope for action exploration.

Comparing ClothPPO to PPO Trained from Scratch.

Fig.5 provides a comparative visualization of ClothPPO and PPO, the latter of which has been trained from scratch without pre-training. The red line, representing ClothPPO, consistently surpasses the green line, symbolizing PPO From Scratch. ClothPPO not only commences with a higher initial coverage but also exhibits an evident upward progression, indicating the productive impact of pretraining on the learning process. Conversely, the PPO From Scratch method struggles to exhibit significant advancements throughout the training. These findings suggest that employing pretraining with PPO offers a remarkable advantage for learning efficiency and end performance in challenging cloth unfolding tasks when compared to starting from scratch. Pretraining gives the model a head start and leads to more robust learning.

Reward Ablation. As Fig. 4, it is evident that different rewards influence the learning curve and the speed at which the agent achieves high rewards. The loss/entropy graph indicates the variability and uncertainty in the model’s predictions, with lower values suggesting better model stability. Finally, the best coverage graph illustrates the highest coverage

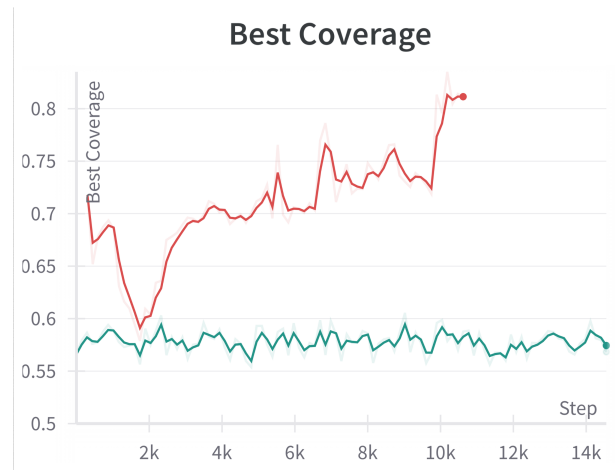


Figure 5: Comparing ClothPPO to PPO From Scratch. The dotted and solid lines represent the original and 0.2-smoothed data respectively. The red line corresponds to ClothPPO, while the blue line represents PPO From Scratch. ClothPPO performance (red line) demonstrates a superior mean best coverage across the number of steps, indicating enhanced task performance.

achieved by the agent, serving as a direct measure of the performance for each reward function. The threshold achievement reward shows a trend towards higher coverage, indicating a potentially more robust learning outcome compared to the other rewards. This could be because the rewards for over achievement have a large variance. This causes the model to struggle more and undergo severe fluctuations. On the other hand, the rewards for immediate termination are excessively challenging, making it difficult for the model to learn.

Reward Scale. As shown in Fig.6, our results indicate that the reward scale can effectively facilitate the convergence of critic networks. The blue line, which represents the scenario with reward normalization, shows a significantly smoother descent and lower variance in critic loss compared

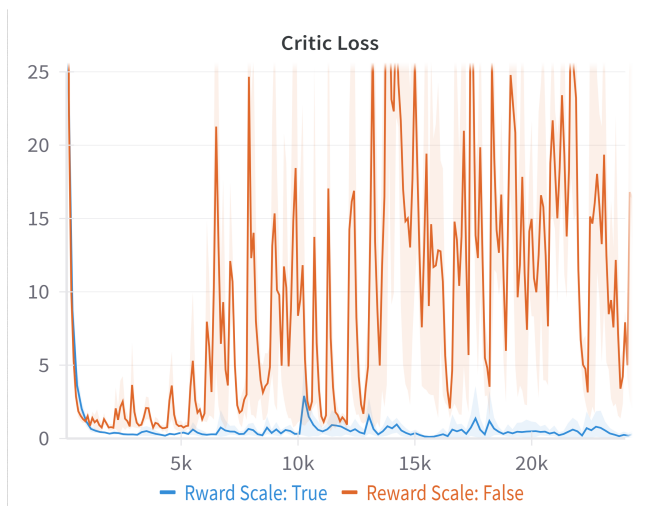


Figure 6: Critic loss comparison. Using the reward scale as reward normalization contributes positively to the learning performance, enhancing stability and efficiency, as evidenced by the critic’s lower and more stable loss values.

to the orange line, which indicates the scenario without reward normalization. Reward normalization leads to a more stable training process, as indicated by the less volatile and more consistent decline in loss. The normalized reward approach can converge faster to a lower loss value than the non-normalized approach. Non-normalized reward settings experience sharp increases in losses, suggesting that wildly changing sparse rewards can produce large prediction errors in the critic model.

7 Conclusion

Our research introduces a novel framework ClothPPO that harnessing the potential of using PPO to enhance the pre-trained model. This approach is aimed to surmount the initial barriers often faced in high-dimensional control tasks and solved the problem of large grasping action space in the task of folding clothes by robots. By optimizing and updating the strategy, the method can increase the robot’s surface area for clothing unfolding under the soft-body manipulation task. Furthermore, our approach is scalable and holds the potential for extension to a broader range of tasks in the future.

Acknowledgments

This work is supported by the National Natural Science Foundation of China (62102152) and sponsored by the Special Fund for Graduate International Conferences of East China Normal University.

References

- [Avigal *et al.*, 2022] Yahav Avigal, Lars Berscheid, Tamim Asfour, Torsten Kröger, and Ken Goldberg. Speedfolding: Learning efficient bimanual folding of garments. In *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1–8. IEEE, 2022.
- [Bertiche *et al.*, 2020] Hugo Bertiche, Meysam Madadi, and Sergio Escalera. Cloth3d: clothed 3d humans. In *European Conference on Computer Vision*, pages 344–359. Springer, 2020.
- [Brockman *et al.*, 2016] Greg Brockman, Vicki Cheung, Ludwig Pettersson, Jonas Schneider, John Schulman, Jie Tang, and Wojciech Zaremba. Openai gym. *arXiv preprint arXiv:1606.01540*, 2016.
- [Canberk *et al.*, 2022] Alper Canberk, Cheng Chi, Huy Ha, Benjamin Burchfiel, Eric Cousineau, Siyuan Feng, and Shuran Song. Cloth funnels: Canonicalized-alignment for multi-purpose garment manipulation. In *International Conference of Robotics and Automation (ICRA)*, 2022.
- [Chen* *et al.*, 2023] Siwei Chen*, Yiqing Xu*, Cunjun Yu*, Linfeng Li, Xiao Ma, Zhongwen Xu, and David Hsu. Daxbench: Benchmarking deformable object manipulation with differentiable physics. In *ICLR*, 2023.
- [Chen *et al.*, 2024] Jiawei Chen, Lin Chen, Jiang Yang, Tianqi Shi, Lechao Cheng, Zunlei Feng, and Mingli Song. Life regression based patch slimming for vision transformers. *Neural Networks*, page 106340, 2024.
- [Doumanoglou *et al.*, 2016] Andreas Doumanoglou, Jan Stria, Georgia Peleka, Ioannis Mariolis, Vladimir Petrik, Andreas Kargakos, Libor Wagner, Václav Hlaváč, Tae-Kyun Kim, and Sotiris Malassiotis. Folding clothes autonomously: A complete pipeline. *IEEE Transactions on Robotics*, 32(6):1461–1478, 2016.
- [François-Lavet *et al.*, 2018] Vincent François-Lavet, Peter Henderson, Riashat Islam, Marc G Bellemare, Joelle Pineau, et al. An introduction to deep reinforcement learning. *Foundations and Trends® in Machine Learning*, 11(3-4):219–354, 2018.
- [Ganapathi *et al.*, 2021] Aditya Ganapathi, Priya Sundaresan, Brijen Thananjeyan, Ashwin Balakrishna, Daniel Seita, Jennifer Grannen, Minh Hwang, Ryan Hoque, Joseph E Gonzalez, Nawid Jamali, et al. Learning dense visual correspondences in simulation to smooth and fold real fabrics. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11515–11522. IEEE, 2021.
- [Gupta *et al.*, 2021] Agrim Gupta, Silvio Savarese, Surya Ganguli, and Li Fei-Fei. Embodied intelligence via learning and evolution. *Nature communications*, 12(1):5721, 2021.
- [Ha and Song, 2021] Huy Ha and Shuran Song. Flingbot: The unreasonable effectiveness of dynamic manipulation for cloth unfolding. In *Conference on Robotic Learning (CoRL)*, 2021.
- [Hu *et al.*, 2023] Kaiwen Hu, Jing Gao, Fangyuan Mao, Xinhui Song, Lechao Cheng, Zunlei Feng, and Mingli Song. Disassembling convolutional segmentation network. *International Journal of Computer Vision*, 131(7):1741–1760, 2023.
- [Huang *et al.*, 2023] Zixuan Huang, Xingyu Lin, and David Held. Self-supervised cloth reconstruction

- via action-conditioned cloth tracking. *arXiv preprint arXiv:2302.09502*, 2023.
- [Li *et al.*, 2018] Yunzhu Li, Jiajun Wu, Russ Tedrake, Joshua B Tenenbaum, and Antonio Torralba. Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids. *arXiv preprint arXiv:1810.01566*, 2018.
- [Lin *et al.*, 2020] Xingyu Lin, Yufei Wang, Jake Olkin, and David Held. Softgym: Benchmarking deep reinforcement learning for deformable object manipulation. In *Conference on Robot Learning*, 2020.
- [Lin *et al.*, 2022] Xingyu Lin, Yufei Wang, Zixuan Huang, and David Held. Learning visible connectivity dynamics for cloth smoothing. In *Conference on Robot Learning*, pages 256–266. PMLR, 2022.
- [Matas *et al.*, 2018] Jan Matas, Stephen James, and Andrew J Davison. Sim-to-real reinforcement learning for deformable object manipulation. In *Conference on Robot Learning*, pages 734–743. PMLR, 2018.
- [Mo *et al.*, 2022] Kai Mo, Chongkun Xia, Xueqian Wang, Yuhong Deng, Xuehai Gao, and Bin Liang. Foldsformer: Learning sequential multi-step cloth manipulation with space-time attention. *IEEE Robotics and Automation Letters*, 8(2):760–767, 2022.
- [Mulero *et al.*, 2023] David Blanco Mulero, Gökhan Alcan, Fares Abu-Dakka, and Ville Kyrki. Qdp: Learning to sequentially optimise quasi-static and dynamic manipulation primitives for robotic cloth manipulation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2023.
- [Nair *et al.*, 2017] 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 *2017 IEEE international conference on robotics and automation (ICRA)*, pages 2146–2153. IEEE, 2017.
- [Ouyang *et al.*, 2022] Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- [Ramamurthy *et al.*, 2022] Rajkumar Ramamurthy, Prithviraj Ammanabrolu, Kianté Brantley, Jack Hessel, Rafet Sifa, Christian Bauckhage, Hannaneh Hajishirzi, and Yejin Choi. Is reinforcement learning (not) for natural language processing?: Benchmarks, baselines, and building blocks for natural language policy optimization. 2022.
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [Seita *et al.*, 2021] Daniel Seita, Pete Florence, Jonathan Tompson, Erwin Coumans, Vikas Sindhwani, Ken Goldberg, and Andy Zeng. Learning to rearrange deformable cables, fabrics, and bags with goal-conditioned transporter networks. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4568–4575. IEEE, 2021.
- [Sondik, 1971] Edward Jay Sondik. *The optimal control of partially observable Markov processes*. Stanford University, 1971.
- [Triantafyllou *et al.*, 2016] Dimitra Triantafyllou, Ioannis Mariolis, Andreas Kargakos, Sotiris Malassiotis, and Nikos Aspragathos. A geometric approach to robotic unfolding of garments. *Robotics and Autonomous Systems*, 75:233–243, 2016.
- [Tutz, 1991] Gerhard Tutz. Sequential models in categorical regression. *Computational Statistics & Data Analysis*, 11(3):275–295, 1991.
- [Weng *et al.*, 2022] Jiayi Weng, Huayu Chen, Dong Yan, Kaichao You, Alexis Duburcq, Minghao Zhang, Yi Su, Hang Su, and Jun Zhu. Tianshou: A highly modularized deep reinforcement learning library. *Journal of Machine Learning Research*, 23(267):1–6, 2022.
- [Wu *et al.*, 2020a] Jimmy Wu, Xingyuan Sun, Andy Zeng, Shuran Song, Johnny Lee, Szymon Rusinkiewicz, and Thomas Funkhouser. Spatial action maps for mobile manipulation. *arXiv preprint arXiv:2004.09141*, 2020.
- [Wu *et al.*, 2020b] Yilin Wu, Wilson Yan, Thanard Kurutach, Lerrel Pinto, and Pieter Abbeel. Learning to manipulate deformable objects without demonstrations, 2020.
- [Xiong *et al.*, 2023] Haoyu Xiong, Haoyuan Fu, Jieyi Zhang, Chen Bao, Qiang Zhang, Yongxi Huang, Wenqiang Xu, Animesh Garg, and Cewu Lu. Robotube: Learning household manipulation from human videos with simulated twin environments. In *Conference on Robot Learning*, pages 1–10. PMLR, 2023.
- [Zeng *et al.*, 2020] Andy Zeng, Pete Florence, Jonathan Tompson, Stefan Welker, Jonathan Chien, Maria Attarian, Travis Armstrong, Ivan Krasin, Dan Duong, Vikas Sindhwani, and Johnny Lee. Transporter networks: Rearranging the visual world for robotic manipulation. *Conference on Robot Learning (CoRL)*, 2020.