

# Learning Language-Conditioned Deformable Object Manipulation with Graph Dynamics

Yuhong Deng<sup>1,†</sup>, Kai Mo<sup>2,†</sup>, Chongkun Xia<sup>2</sup> and Xueqian Wang<sup>2,\*</sup>

**Abstract**—Multi-task learning of deformable object manipulation is a challenging problem in robot manipulation. Most previous works address this problem in a goal-conditioned way and adapt goal images to specify different tasks, which limits the multi-task learning performance and can not generalize to new tasks. Thus, we adapt language instruction to specify deformable object manipulation tasks and propose a learning framework. We first design a unified Transformer-based architecture to understand multi-modal data and output picking and placing action. Besides, we have applied the visible connectivity graph to tackle nonlinear dynamics and complex configuration of the deformable object. Both simulated and real experiments have demonstrated that the proposed method is effective and can generalize to unseen instructions and tasks. Compared with the state-of-the-art method, our method achieves higher success rates (87.2% on average) and has a 75.6% shorter inference time. We also demonstrate that our method performs well in real-world experiments. Supplementary videos can be found at <https://sites.google.com/view/language-deformable>.

## I. INTRODUCTION

Recently, vision-based deformable object manipulation has been widely investigated. The robot is supposed to infer a sequence of manipulation actions from visual observations to manipulate a deformable object into a prescribed goal configuration [1], [2]. Early works focus on learning task-specific manipulation skills. A generic robot that can complete different deformable object manipulation tasks is still an open challenge. Toward this goal, some researchers have made progress in using goal images to specify different tasks and achieve multi-task learning [3], [4]. However, goal images often over-defined tasks by information not related to the task requirement, such as position, color, and image style. Besides, the model can not deal with unseen goal images from different domains and generalize to new tasks.

Compared with goal image, language instruction only provide necessary information related to task requirement and can define tasks flexibly. The prior knowledge of the pre-trained language model can also help the robot deal with unseen tasks. Thus, language-conditioned manipulation has been a new trend to pave the way for the generic robot that can complete various manipulation tasks. However, previous language-conditioned manipulation learning methods

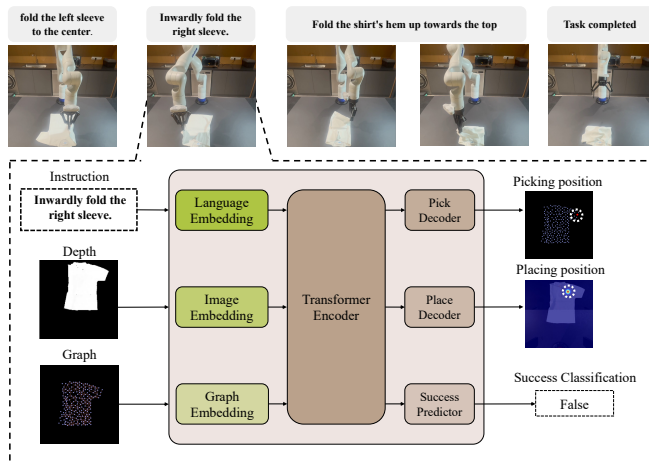


Fig. 1. **Overview.** We design a unified Transformer-based model and apply graph representation to solve language-conditioned deformable object manipulation tasks. Our model performs well on deformable object manipulation tasks.

are limited to templated manipulations like picking, placing, and moving rigid objects [5], [6]. Language-conditioned deformable object manipulation tasks have not been solved for two reasons. Deformable object manipulation tasks are often sequential multi-step tasks, which pose higher requirements for inference and decision-making at the intersection of vision, language, and action. In addition, the manipulation task is more difficult because the nonlinear dynamics [7] and complex configurations [8] of deformable objects.

We design a unified Transformer-based model architecture to tackle the challenge of inference and decision-making on multi-modal data (Fig. 1). We first implement a CLIP [9] model pre-trained on millions of image-caption pairs to provide language embeddings, which can provide a powerful prior for grounding language in visual scenes. As for the image, we borrow the method in ViT [10] (Vision Transformer) and get the visual embeddings by linear projection. To deal with deformable dynamics and configurations, we establish a visible connectivity graph [11] and get the graph embeddings by linearly projecting representation vectors of graph nodes. The visible connectivity graph can represent the deformable object’s spatial structure and overcome the challenges of partial observability and self-occlusions. After obtaining the language, image, and graph embeddings, all embeddings will pass through the type embedding layer [12] to achieve multi-modal data fusion. Finally, a Transformer [13] encoder-decoder structure is used to generate picking and placing action possibility distribution, where positions with max possibility will be picking and placing positions.

<sup>†</sup> indicates the authors with equal contributions.

<sup>1</sup> National University of Singapore, Singapore

<sup>2</sup> Tsinghua Shenzhen International Graduate School, Shenzhen, China

\*Correspondence to: wang.xq@sz.tsinghua.edu.cn

This work was supported by the National Key R&D Program of China (2022YFB4701400/4701402), National Natural Science Foundation of China (No. U21B6002, 62203260, 92248304), The Shenzhen Science Fund for Distinguished Young Scholars (RCJC20210706091946001), Guangdong Basic and Applied Basic Research Foundation (2023A1515011773).

Since there is no existing dataset for language-conditioned deformable object manipulation tasks, we build a dataset including various language instructions and corresponding deformable object manipulation tasks in the simulation environment. We then conduct simulated experiments to evaluate our proposed method. The results demonstrate that our method outperforms the state-of-the-art framework [5] on multi-task learning of deformable object manipulation. Our model can also generalize to unseen instructions and tasks not provided in the training process. Real-world experiments also demonstrate that our model trained in simulation can be zero-shot transferred to the real world. The contributions of this work are summarized as follows:

- 1) We propose a novel learning framework that extends instruction-following robots’ application on sequential multi-step deformable object manipulation;
- 2) We propose a unified Transformer-based model architecture for language-conditioned deformable object manipulation policy learning;
- 3) We conduct simulated and real-world experiments to demonstrate that the proposed framework outperforms the state-of-the-art method and can generalize to unseen instructions and tasks.

The rest of this paper is organized as follows. The related work is reviewed in Section II. Section III presents details of the learning framework design. The experimental setup and results are provided in Section IV. Section V concludes this paper and discusses future works.

## II. RELATED WORK

### A. Learning for Deformable Object Manipulation

Learning-based methods have been widely used to equip the robot with advanced deformable object manipulation abilities. There are two main learning-based methods: Model-based methods rely on a forward dynamics model, which can predict the configurations of deformable objects under a given action. The critical issue is to obtain accurate dynamics. Learning accurate deformable dynamics from interaction data between robots and objects has become a solution [14], [15]. However, the generalization of the learning model on different objects and tasks is still an open challenge. Policy-based methods learn manipulation policy directly from observation without a forward dynamics model. The robot learns manipulation policy from expert demonstrations [16] or exploratory robot interactions [17]. However, most previous policy-based methods are limited to task-specific policy, which is inefficient in real-world applications [18]. Some approaches learn a goal-conditioned policy to provide a general framework for deformable object manipulation, where the goal image specifies the task [3], [4]. Unlike previous work, we adapt language instructions to specify different deformable object manipulation tasks and achieve a general learning framework.

### B. Language-conditioned Robotic Manipulation policy

Language-conditioned robotic manipulation has received growing attention recently, where the manipulation task

is specified by language instruction [19], [20]. Language instruction specification is much easier to obtain than goal image specification and can specify more diversified tasks feasibly and effectively. Despite the abundant benefits of commanding robots with natural language, such agents require deep integration of multiple data modalities (language, vision, action). Language-Conditioned Imitation Learning addresses this problem by mapping actions directly from vision and language understanding in an end-to-end fashion [5]. However, collecting demonstrations paired with language annotations in the real robot is costly and time-consuming. To tackle this data-collecting problem, Nair et al. [21] adopt crowd-sourced annotation to obtain sufficient datasets. Although some works improve the model’s performance in more diverse languages [22], [23], previous language-conditioned robotic manipulation approaches are limited in rigid objects. The robot only needs to manipulate a rigid object with low-level skills such as grasping and pushing. Simplifying objects and actions will limit the application scenarios. Thus, we propose a solution framework that can be applied to complex deformable object manipulation tasks.

## III. METHODS

### A. Problem Formulation

Our goal is to learn a language-conditioned deformable object manipulation policy  $\pi_{\theta}$  parameterized by  $\theta$ , which can generate a sequence of manipulation action  $\{\mathbf{a}_t\} (t = 0, 1, 2, \dots, T)$  in a closed-loop manner from a language instruction  $\mathbf{s}$  and the current visual observation  $\mathbf{o}_t$ :

$$\begin{aligned} \mathbf{a}_t &\leftarrow \pi_{\theta}(\mathbf{o}_t, \mathbf{s}) \\ \text{and } \mathbf{o}_{t+1} &\leftarrow \mathcal{T}(\mathbf{o}_t, \mathbf{a}_t) \end{aligned} \quad (1)$$

where  $\mathcal{T}$  denotes the state transition. Since the deformable object has a complex configuration, the visual observation  $\mathbf{o}$  is composed of two parts in our solution framework: the top-down depth image  $\mathbf{I}$  and the representation graph  $\mathbf{G}$ :

$$\mathbf{a}_t \leftarrow \pi_{\theta}(\mathbf{I}_t, \mathbf{G}_t, \mathbf{s}) \quad (2)$$

The action space is defined as picking and placing action:

$$\mathbf{a}_t = \{\mathbf{a}_t^{\text{pick}}, \mathbf{a}_t^{\text{place}}\} \quad (3)$$

where  $\mathbf{a}_t^{\text{pick}}$  and  $\mathbf{a}_t^{\text{place}}$  denote the robot end-effector’s picking and placing poses.

We formulate the problem as a supervised learning problem. The policy  $\pi_{\theta}$  is learned from expert demonstrations.

### B. Model Architecture

We design a unified transformer-based model architecture (Fig. 2) for learning language-conditioned deformable object manipulation policy. We first leverage different modules to embed these multi-modal data. Then an encoder-decoder structure is used to generate picking and placing actions from multi-modal embeddings.

**Language Embedding:** We employ the language encoder in the CLIP model [9] to embed the language instruction  $\mathbf{s}$ , which can provide a powerful prior for grounding language in visual scenes. Like previous works [10], [24], we prepend

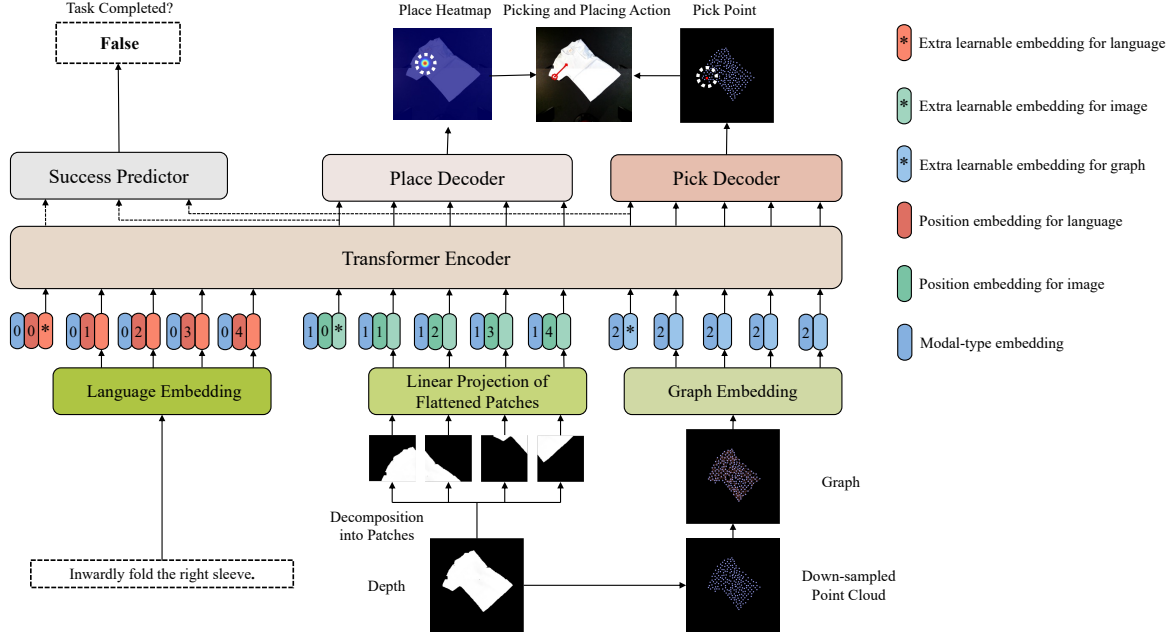


Fig. 2. **Method overview.** We design a unified Transformer-based model architecture to understand the multi-modal data and output picking and placing action with task completion prediction. We apply a visible connectivity graph to tackle deformable objects’ complex configurations and dynamics.

an extra learnable embedding for aggregating the language representation and add learnable position embeddings to retain positional information. Mathematically, the language embedding  $z_s$  is computed as follows:

$$z_s = [S_{head}; \text{CLIP}(s)] + S_{pos} \quad (4)$$

where  $\text{CLIP}()$  denotes the CLIP’s language encoder,  $S_{head}$  denotes the extra learnable embedding,  $S_{pos}$  denotes the position embeddings.

**Image Embedding:** Inspired by the previous work [25], we use depth images rather than RGB images to make our framework transferable to the real world without any additional training in the real world. We first decompose the depth image  $\mathbf{I} \in \mathbb{R}^{H \times W \times 1}$  into  $N$  non-overlapping patches. We then flatten these 2D patches into 1D vectors  $x_I \in \mathbb{R}^{H \times W/N}$ . We linearly map these vectors to align the CLIP output’s dimension. We prepend an extra learnable embedding to aggregate the image representation and add learnable position embeddings for retaining positional information. Mathematically, The image embedding  $z_I$  is computed as follows:

$$z_I = [I_{head}; W_I x_I^1; W_I x_I^2; \dots, W_I x_I^N] + I_{pos} \quad (5)$$

where  $W_I$  denotes a learnable matrix,  $I_{head}$  denotes the extra learnable embedding,  $I_{pos}$  denotes the position embeddings.

**Graph Embedding:** We apply a visible connectivity graph  $\langle V, E \rangle$  to pose the inductive bias of the deformable objects’ physics to the framework [11]. The nodes  $V$  represent the particles that compose the deformable object. Specifically, the down-sampled point cloud  $\mathbf{P} = \{v_i\}_{i=1, \dots, K}$  of the deformable object constructs the nodes  $V$ . There are two types of edges  $E$ , nearby edges  $E^C$  model the collision between nodes and mesh edges  $E^M$  model the deformable object’s spatial structure. We construct nearby edges  $E^C$

based on the euclidean distance of every two nodes firstly:

$$E^C = \{e_{ij} | \|v_i - v_j\|_2 < R\} \quad (6)$$

where  $R$  is a distance threshold, and  $v_i, v_j$  are the positions of the nodes  $i, j$ .

We then use the pre-trained edge GNN [11] to predict the mesh edges  $E^M$  from the graph  $\langle \mathbf{P}, E^C \rangle$ . We also prepend an extra learnable embedding that can aggregate the graph representation. We do not add position embeddings because Transformer without position embeddings is permutation-invariant and a natural fit for graphs [26]. Mathematically, The graph embedding  $z_G$  is computed as follows:

$$z_G = [G_{head}; G_{edge}(\langle \mathbf{P}, E^C \rangle)] \quad (7)$$

where  $G_{head}$  denotes the extra learnable embedding,  $G_{edge}()$  denotes the pre-trained edge GNN.

**Encoder:** After obtaining the multi-modal embeddings, we make an aggregation by concatenating these multi-modal embeddings and adding modal-type embeddings:

$$z_0 = [z_s + S_{type}; z_I + I_{type}; z_G + G_{type}] \quad (8)$$

where  $S_{type}, I_{type}, G_{type}$  are modal-type embeddings of the language, the image and the graph, respectively.

We then input the vector  $z_0$  to a Transformer encoder proposed in ViT [10], which has  $L$  layers:

$$\begin{aligned} z'_l &= \text{MSA}(\text{LN}(z_{l-1})) + z_{l-1} & l = 1, \dots, L \\ z_l &= \text{MLP}(\text{LN}(z'_l)) + z'_l & l = 1, \dots, L \end{aligned} \quad (9)$$

where  $\text{MSA}()$  denotes multiheaded self-attention,  $\text{LN}()$  denotes LayerNorm, and  $\text{MLP}()$  denotes an MLP layer.

**Pick Decoder and Place Decoder:** We then design a pick decoder and a place decoder to predict picking and placing position from the output embeddings of the encoder layer. The pick decoder is an MLP containing one layer, while

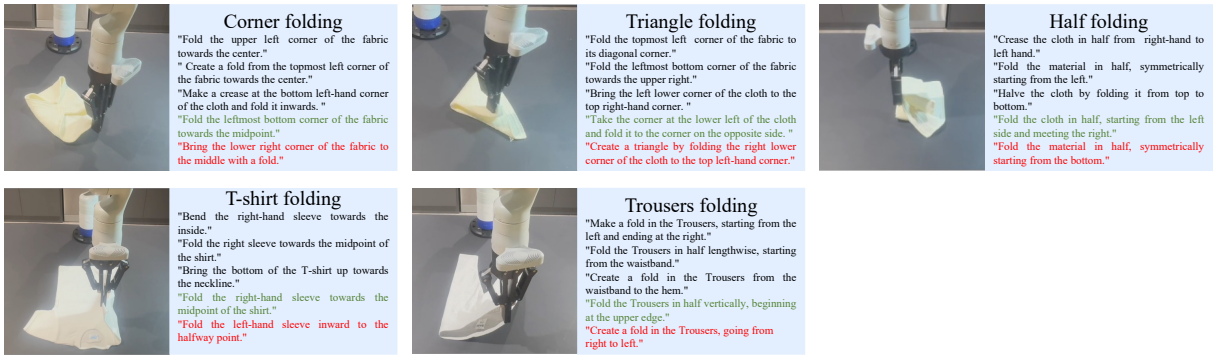


Fig. 3. **Some examples of language-conditioned deformable object manipulation Tasks.** Seen instructions, unseen instructions, unseen tasks are marked in black, grey and red, respectively.

the place decoder consists of convolutional and upsampling layers alternately. The pick decoder takes the graph’s node embeddings as input and outputs the probability distribution  $Q_{\text{pick}} \in \mathbb{R}^K$ , which represent the probability of each point in the down-sampled point cloud  $\mathbf{P}$  being a picking position. The place encoder takes the image’s patch embeddings as input and outputs a pixel-wise heatmap  $Q_{\text{place}} \in \mathbb{R}^{H \times W}$ , which represent the probability of each pixel in the image being a placing position. The optimal picking action  $\mathbf{a}_{\text{pick}}$  and the optimal placing action  $\mathbf{a}_{\text{place}}$  are computed as follows:

$$\begin{aligned} \mathbf{a}_{\text{pick}} &= \operatorname{argmax}_{\mathbf{a}} Q_{\text{pick}}(\mathbf{a}) \\ \mathbf{a}_{\text{place}} &= \operatorname{argmax}_{\mathbf{a}} Q_{\text{place}}(\mathbf{a}) \end{aligned} \quad (10)$$

**Success Classifier:** Most previous language-conditioned manipulation learning algorithm can not estimate whether a task has been completed [5], [20], [22], which is not practical in real-world employment. Thus, we design a successful classifier to estimate task completion, making our framework more autonomous. The success classifier is an MLP that contains two layers, which can perform binary classification to indicate task completion. The success classifier’s input is the concatenation of the output embeddings corresponding to three extra embeddings  $S_{\text{head}}$ ,  $I_{\text{head}}$  and  $G_{\text{head}}$ .

### C. Implementation Details:

We use behavioral cloning to train a multi-task model. The training data is collected in the SoftGym suite [27]. The 3D models of deformable objects are sampled from CLOTH3D dataset [28]. In SoftGym, the deformable objects are modeled as particles whose ground truth positions and velocities can be accessed. Thus, we can easily collect expert demonstrations using a oracle demonstrator. The object’s initial configurations (size, pose, etc.) are randomized during data collection. Each expert demonstration  $\zeta$  is composed of the language instruction  $\mathbf{s}$ , the observation-action pair  $(\mathbf{o}_t, \mathbf{a}_t)$  with  $t = 1, 2, \dots, T$ :

$$\zeta = \{l, (\mathbf{o}_1, \mathbf{a}_1), (\mathbf{o}_2, \mathbf{a}_2), \dots, (\mathbf{o}_T, \mathbf{a}_T)\} \quad (11)$$

For model training, we train the success classifier and other modules separately. We first freeze the CLIP encoder and the edge GNN and then train the modules except for the classifier with the binary cross-entropy (BCE) loss between the

predicted  $Q_{\text{pick}}, Q_{\text{place}}$  and the ground truth  $Q_{\text{pick}}^{\text{gt}}, Q_{\text{place}}^{\text{gt}}$ :

$$\mathcal{L}_{\text{action}} = \text{BCE}(Q_{\text{pick}}, Q_{\text{pick}}^{\text{gt}}) + \text{BCE}(Q_{\text{place}}, Q_{\text{place}}^{\text{gt}}) \quad (12)$$

Finally, We freeze the trained modules and train the success classifier with the BCE loss.

## IV. EXPERIMENTS

This section presents simulated and real-world experiments to answer the following questions: 1) How well does our framework perform on multi-task learning of deformable object manipulation compared with the baseline methods? 2) What role does the visible connectivity graph play in our algorithm design? 3) How well does our framework generalize to unseen instructions and tasks? and 4) How well does our framework perform on real-world language-conditioned deformable object manipulation tasks?

### A. Simulation Experiments Setup

All simulation experiments are conducted in the SoftGym suite [27]. The robot is provided with language instruction and is supposed to complete the instruction with only current visual observation inputs. We evaluate the performance of our framework with 5 types of language-conditioned manipulation tasks. Fig 3 shows some examples, and TABLE II shows the number of instructions associated with each type of task. There are various instructions and tasks in the simulation experiments. The testing instructions can be divided into three parts: seen instructions, unseen instructions, and unseen tasks. Seen instructions means instructions seen in the training process. The manipulation tasks that unseen instructions and seen instructions specify are the same. The difference is the language description. Unseen instructions can not be seen in the training process. Unseen tasks mean new tasks not seen in training. Unseen tasks require new manipulation skills that the model has not learned in training. For example, for corner folding, folding from the bottom left, top right, and top left can be seen in training; folding from the bottom right is an unseen task. For T-shirt folding, folding the right sleeve can be seen; folding the left sleeve is an unseen task.

TABLE I

**SIMULATION EXPERIMENT RESULTS.** THE AVERAGE SUCCESS RATES (%) ON TESTING TASKS (SEEN INSTRUCTIONS, UNSEEN INSTRUCTIONS AND UNSEEN TASKS). MODELS ARE TRAINED WITH 100 AND 1000 DEMONSTRATIONS PER TASK. THE BEST PERFORMANCE IS IN BOLD.

Method	corner folding (seen instructions)		triangle folding (seen instructions)		half folding (seen instructions)		T-shirt folding (seen instructions)		Trousers folding (seen instructions)	
	100	1000	100	1000	100	1000	100	1000	100	1000
Foldsformer [25]	80.0	90.0	52.0	68.0	16.0	26.0	18.0	24.0	8.0	20.0
CLIPORT [5]	78.0	86.0	<b>76.0</b>	78.0	50.0	56.0	54.0	74.0	38.0	46.0
ours (w/o graph)	100.0	100.0	66.0	84.0	44.0	56.0	78.0	78.0	64.0	<b>88.0</b>
ours (full method)	<b>100.0</b>	<b>100.0</b>	72.0	<b>92.0</b>	<b>52.0</b>	<b>74.0</b>	<b>86.0</b>	<b>84.0</b>	<b>74.0</b>	86.0

Method	corner folding (unseen instructions)		triangle folding (unseen instructions)		half folding (unseen instructions)		T-shirt folding (unseen instructions)		Trousers folding (unseen instructions)	
	100	1000	100	1000	100	1000	100	1000	100	1000
CLIPORT [5]	80.0	90.0	76.0	74.0	38.0	50.0	56.0	70.0	32.0	38.0
ours (w/o graph)	<b>100.0</b>	100.0	72.0	78.0	36.0	60.0	<b>86.0</b>	80.0	54.0	<b>84.0</b>
ours (full method)	96.0	<b>100.0</b>	<b>80.0</b>	<b>92.0</b>	<b>40.0</b>	<b>68.0</b>	80.0	<b>80.0</b>	<b>74.0</b>	<b>92.0</b>

Method	corner folding (unseen tasks)		triangle folding (unseen tasks)		half folding (unseen tasks)		T-shirt folding (unseen tasks)		Trousers folding (unseen tasks)	
	100	1000	100	1000	100	1000	100	1000	100	1000
CLIPORT [5]	70.0	76.0	70.0	74.0	0.0	0.0	22.0	46.0	0.0	8.0
ours (w/o graph)	<b>74.0</b>	<b>80.0</b>	50.0	80.0	0.0	<b>2.0</b>	28.0	76.0	<b>10.0</b>	14.0
ours (full method)	36.0	78.0	<b>76.0</b>	<b>88.0</b>	<b>0.0</b>	0.0	<b>42.0</b>	<b>80.0</b>	0.0	<b>32.0</b>

TABLE II

**TESTING INSTRUCTIONS IN SIMULATION EXPERIMENTS.**

Task	seen instructions	unseen instructions	unseen tasks
corner folding	192	48	64
triangle folding	192	48	64
half folding	192	48	64
T-shirt folding	96	24	48
Trousers folding	96	24	48
Total	768	192	288

### B. Simulation Experiment Results

We compare the performance of our methods with two baseline methods: Foldsformer [25] is a state-of-the-art method that adapts a sequence of sub-goal images to specify the task. Foldsformer is provided with sub-goal images rather than language instruction in each task. CLIPORT [5] represents the typical network for language-conditioned manipulation learning. It relies on a two-stream architecture and uses pre-trained vision-language models for language-conditioned manipulation policies. Besides, we set an ablation study to evaluate the role of the graph in our framework. Ours (w/o graph) has the same backbone architecture as the proposed algorithm. The only difference is that we did not provide it with graph information.

We first evaluate the success rate on seen instructions, unseen instructions, and unseen tasks (each type of task has 50 task instances). We train all models with 100 and 1000 demonstrations separately. The success metric is the mean particle position error between the cloth states achieved by the policy and an oracle demonstrator. We define a task as a success if the mean particle position error is less than 0.0125m (the diameter of a particle in SoftGym). The results are shown in TABLE I.

Overall, Foldsformer performs worst, illustrating the mer-

TABLE III

**MODEL CAPACITY.** WE COMPARE MODEL PARAMETERS, FLOPS, AND INFERENCE TIME OF DIFFERENT MODELS. THE BEST PERFORMANCE IS IN BOLD.

Method	Params(M)	FLOPS(G)	Inference Time (ms)
Foldsformer [25]	<b>11.6</b>	24.58	19
CLIPORT [5]	423.19	384.37	131
ours (w/o graph)	51.03	<b>19.81</b>	<b>15</b>
ours (full method)	53.55	34.38	32

its of language specification in multi-task learning of robot manipulation. Compared with a demonstration of an image sequence, natural language can provide sufficient cues of task requirements without over-defining the task by the object’s texture, position, and size. Besides, a demonstration of an image sequence could only capture one instance of success [29].

Ours (w/o graph) generally outperforms CLIPORT, demonstrating that the proposed model architecture effectively learns language-conditioned manipulation policy. We use a unified Transformer-based model architecture to deal with multi-modal input and output picking and placing action. Such model architecture outperforms the state-of-the-art method.

By introducing a visible connectivity graph, ours (full method) outperforms all baseline methods, especially on tasks with more steps and complex deformable objects (such as T-shirt and trousers folding). In these more complicated tasks, the manipulation process is more likely to cause irregular self-occlusion and partial observation. The visible connectivity graph can help the robot accurately capture deformable dynamics and configurations.

In addition, our model is also more data-efficient. In TABLE I, the success rates of our model trained on 100

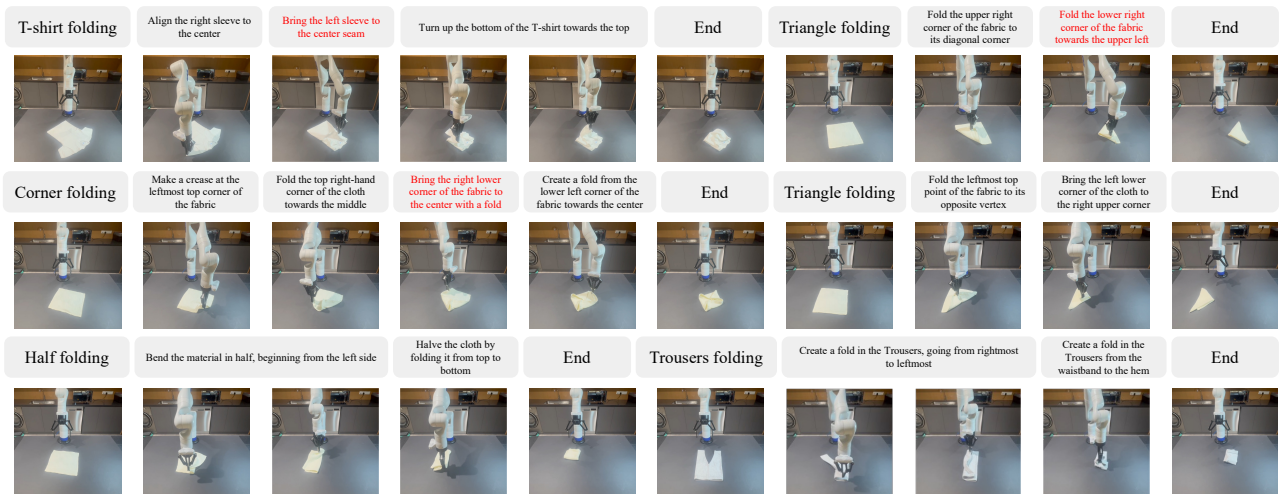


Fig. 4. **Real World Experiments.** Our model performs well in language-conditioned deformable object manipulation tasks and can generalize to unseen tasks in the real world. Unseen tasks are marked in red.

demonstrations are higher than those of baseline models trained on 1000 demonstrations on almost all involved tasks.

The experiment results also show that our model can be general to unseen instructions and even unseen tasks. Our model can understand unseen instructions with the prior knowledge of the pre-trained language model. More importantly, our model learns to ground the spatial displacement hidden in the instruction in the visual image instead of just memorizing some picking and placing positions, which makes the model general to unseen tasks such as folding from different directions and folding different parts. We also found that it is hard for the model to generalize to unseen tasks such as half folding, where the manipulation is more likely to lead to an arbitrary state. In these tasks, the model must see enough arbitrary visual features in the training.

Besides, our model is much lighter than the baseline method. We calculate different models’ FLOPs (Floating Point Operations), model parameters, and inference times. The results are shown in TABLE. III. Overall, Foldsformer is the lightest because it does not deal with multi-modal data. Ours (full method) and ours (w/o graph) are dramatically lighter in model FLOPs, parameters, and inference time than CLIPORT, which is attributed to the proposed unified transformer-based architecture. It can also be seen that improving model performance by introducing graph representation is feasible in terms of model efficiency.

### C. Real World Experiments

As described previously, we use depth images rather than RGB images to make our framework able to be transferred to the real world directly. We evaluate the sim-to-real performance of our framework on a kinova robot with a standard two-finger Robotiq gripper. A Realsense RGB-D camera mounted on the robot’s end-effector is used to capture visual observations. The deformable object (towel, T-shirt, or trousers) is placed on the platform in front of the robot. We evaluate the success rate of our model on 5 types of real-world language-conditioned manipulation tasks. We define a task as success according to the Mean Intersection Union

TABLE IV

**REAL EXPERIMENT RESULTS.** THE AVERAGE SUCCESS RATE (%) ON TESTING TASKS.

Task	success rate(%)
corner folding	80.0
triangle folding	70.0
half folding	40.0
T-shirt folding	60.0
Trousers folding	40.0

(MIoU) between the deformable object masks achieved by our model and the human experts. If the MIoU exceeds 0.9, the task is successful. Table. IV shows the results. Real experiments’ success rate is close to simulation experiments’ success rate.

Fig. 4 shows some examples in our real-world experiments. Our model can complete both seen and unseen tasks in real-world experiments. A complete video recording of these experiments can be found on our project website.

### V. CONCLUSION

In this paper, we propose a novel framework for language-conditioned deformable manipulation policy learning. We design a unified Transformer-based model architecture to deal with multi-modal data and output precise picking and placing action. Besides, we construct a visible connective graph to represent the spatial structure of the deformable object. Extensive experiments have been conducted to verify the proposed model architecture’s performance, proving that our framework can improve the multi-task learning performance on deformable object manipulation and generalize to unseen instructions and unseen tasks. Furthermore, our framework performs well on real-world experiments. For future work, we will explore the method to deal with arbitrary configurations during manipulation and improve the proposed framework’s performance on more complicated unseen tasks.

## REFERENCES

- [1] B. Thach, B. Y. Cho, A. Kuntz, and T. Hermans, "learning visual shape control of novel 3d deformable objects from partial-view point clouds," in *2022 International Conference on Robotics and Automation (ICRA)*, pp. 8274–8281, IEEE, 2022.
- [2] T. Weng, S. M. Bajracharya, Y. Wang, K. Agrawal, and D. Held, "Fabricflownet: Bimanual cloth manipulation with a flow-based policy," in *Conference on Robot Learning*, pp. 192–202, PMLR, 2022.
- [3] D. Seita, P. Florence, J. Tompson, E. Coumans, V. Sindhwani, K. Goldberg, and A. Zeng, "Learning to rearrange deformable cables, fabrics, and bags with goal-conditioned transporter networks," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4568–4575, IEEE, 2021.
- [4] R. Lee, D. Ward, V. Dasagi, A. Cosgun, J. Leitner, and P. Corke, "Learning arbitrary-goal fabric folding with one hour of real robot experience," in *Conference on Robot Learning*, pp. 2317–2327, PMLR, 2021.
- [5] M. Shridhar, L. Manuelli, and D. Fox, "Cliport: What and where pathways for robotic manipulation," in *Conference on Robot Learning*, pp. 894–906, PMLR, 2022.
- [6] A. Brohan, Y. Chebotar, C. Finn, K. Hausman, A. Herzog, D. Ho, J. Ibarz, A. Irpan, E. Jang, R. Julian, *et al.*, "Do as i can, not as i say: Grounding language in robotic affordances," in *6th Annual Conference on Robot Learning*, 2022.
- [7] S. Zimmermann, R. Poranne, and S. Coros, "Dynamic manipulation of deformable objects with implicit integration," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 4209–4216, 2021.
- [8] J. Zhu, A. Cherubini, C. Dune, D. Navarro-Alarcon, F. Alambeigi, D. Berenson, F. Ficuciello, K. Harada, J. Kober, X. Li, *et al.*, "Challenges and outlook in robotic manipulation of deformable objects," *IEEE Robotics & Automation Magazine*, vol. 29, no. 3, pp. 67–77, 2022.
- [9] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, *et al.*, "Learning transferable visual models from natural language supervision," in *International conference on machine learning*, pp. 8748–8763, PMLR, 2021.
- [10] 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.
- [11] X. Lin, Y. Wang, Z. Huang, and D. Held, "Learning visible connectivity dynamics for cloth smoothing," in *Conference on Robot Learning*, pp. 256–266, PMLR, 2022.
- [12] W. Kim, B. Son, and I. Kim, "Vilt: Vision-and-language transformer without convolution or region supervision," in *International Conference on Machine Learning*, pp. 5583–5594, PMLR, 2021.
- [13] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [14] W. Yan, A. Vangipuram, P. Abbeel, and L. Pinto, "Learning predictive representations for deformable objects using contrastive estimation," in *Conference on Robot Learning*, pp. 564–574, PMLR, 2021.
- [15] X. Ma, D. Hsu, and W. S. Lee, "Learning latent graph dynamics for visual manipulation of deformable objects," in *2022 International Conference on Robotics and Automation (ICRA)*, pp. 8266–8273, IEEE, 2022.
- [16] A. Nair, D. Chen, P. Agrawal, P. Isola, P. Abbeel, J. Malik, and S. Levine, "Combining self-supervised learning and imitation for vision-based rope manipulation," in *2017 IEEE international conference on robotics and automation (ICRA)*, pp. 2146–2153, IEEE, 2017.
- [17] J. Matas, S. James, and A. J. Davison, "Sim-to-real reinforcement learning for deformable object manipulation," in *Conference on Robot Learning*, pp. 734–743, PMLR, 2018.
- [18] Y. Wu, W. Yan, T. Kurutach, L. Pinto, and P. Abbeel, "Learning to manipulate deformable objects without demonstrations," in *Robotics: Science and Systems*, 2020.
- [19] E. Stengel-Eskin, A. Hundt, Z. He, A. Murali, N. Gopalan, M. Gombolay, and G. Hager, "Guiding multi-step rearrangement tasks with natural language instructions," in *Conference on Robot Learning*, pp. 1486–1501, PMLR, 2022.
- [20] M. Shridhar, L. Manuelli, and D. Fox, "Perceiver-actor: A multi-task transformer for robotic manipulation," *arXiv preprint arXiv:2209.05451*, 2022.
- [21] S. Nair, E. Mitchell, K. Chen, S. Savarese, C. Finn, *et al.*, "Learning language-conditioned robot behavior from offline data and crowd-sourced annotation," in *Conference on Robot Learning*, pp. 1303–1315, PMLR, 2022.
- [22] P.-L. Guhur, S. Chen, R. Garcia, M. Tapaswi, I. Laptev, and C. Schmid, "Instruction-driven history-aware policies for robotic manipulations," *arXiv preprint arXiv:2209.04899*, 2022.
- [23] O. Mees, L. Hermann, E. Rosete-Beas, and W. Burgard, "Calvin: A benchmark for language-conditioned policy learning for long-horizon robot manipulation tasks," *IEEE Robotics and Automation Letters*, 2022.
- [24] W. Kim, B. Son, and I. Kim, "Vilt: Vision-and-language transformer without convolution or region supervision," in *International Conference on Machine Learning*, pp. 5583–5594, PMLR, 2021.
- [25] K. Mo, C. Xia, X. Wang, Y. Deng, X. Gao, and B. Liang, "Foldformer: Learning sequential multi-step cloth manipulation with space-time attention," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 760–767, 2023.
- [26] Z. Wu, P. Jain, M. Wright, A. Mirhoseini, J. E. Gonzalez, and I. Stoica, "Representing long-range context for graph neural networks with global attention," *Advances in Neural Information Processing Systems*, vol. 34, pp. 13266–13279, 2021.
- [27] X. Lin, Y. Wang, J. Olkin, and D. Held, "Softgym: Benchmarking deep reinforcement learning for deformable object manipulation," in *Conference on Robot Learning*, pp. 432–448, PMLR, 2021.
- [28] H. Bertiche, M. Madadi, and S. Escalera, "Cloth3d: clothed 3d humans," in *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16*, pp. 344–359, Springer, 2020.
- [29] S. Nair, E. Mitchell, K. Chen, S. Savarese, C. Finn, *et al.*, "Learning language-conditioned robot behavior from offline data and crowd-sourced annotation," in *Conference on Robot Learning*, pp. 1303–1315, PMLR, 2022.