

Policy Adaptation from Foundation Model Feedback

Yuying Ge^{1*} Annabella Macaluso² Li Erran Li^{3†} Ping Luo¹ Xiaolong Wang²
¹University of Hong Kong ²University of California, San Diego ³AWS AI, Amazon
<https://geyuying.github.io/PAFF>

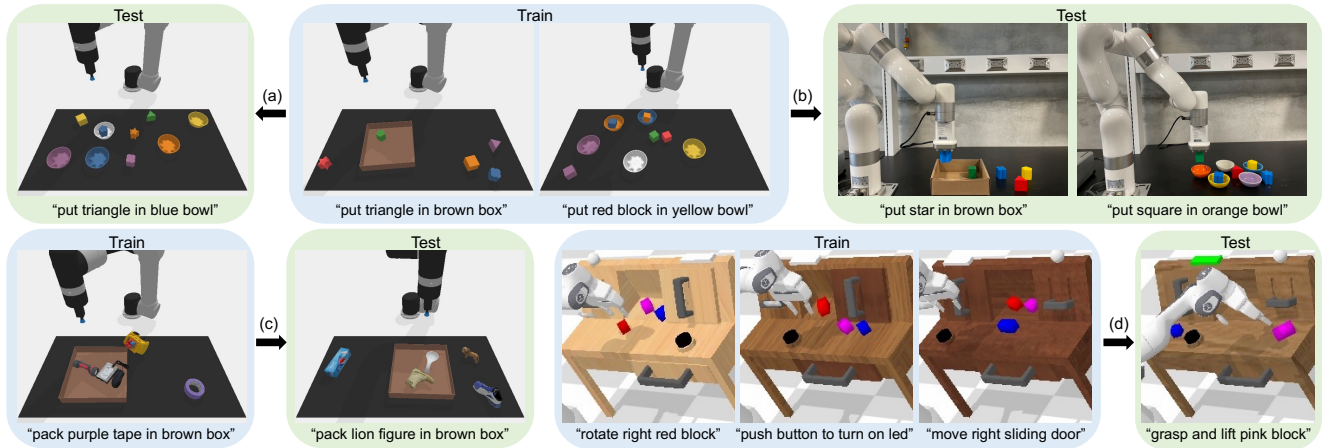


Figure 1. Examples of the language conditioned policy adaptation experiments, including the evaluation of (i) **Compositional Generalization** in (a), where we train a policy to pack *objects of different shapes* in the brown box, and put blocks of different colors in the *bowls of different colors*, and adapt it to put *objects of different shapes* in the *bowls of different colors*; (ii) **Out-of-distribution Generalization** in (c), where we train a policy to pack certain objects in the brown box, and adapt it to pack *unseen objects*, and in (d), where we adapt a policy trained on seen environments to *an unseen environment* with different textures and placements of static elements such as the sliding door and the light button; (iii) **Sim-to-real Transfer** in (b), where we adapt a policy trained on simulation data to the *real world*.

Abstract

Recent progress on vision-language foundation models have brought significant advancement to building general-purpose robots. By using the pre-trained models to encode the scene and instructions as inputs for decision making, the instruction-conditioned policy can generalize across different objects and tasks. While this is encouraging, the policy still fails in most cases given an unseen task or environment. In this work, we propose Policy Adaptation from Foundation model Feedback (PAFF). When deploying the trained policy to a new task or a new environment, we first let the policy play with randomly generated instructions to record the demonstrations. While the execution could be wrong, we can use the pre-trained foundation models to provide feedback to relabel the demonstrations. This automatically provides new pairs of demonstration-instruction data for policy fine-tuning. We evaluate our method on a broad range

of experiments with the focus on generalization on unseen objects, unseen tasks, unseen environments, and sim-to-real transfer. We show PAFF improves baselines by a large margin in all cases.

1. Introduction

Learning generalizable manipulation policies have been a long standing problem in robotics. The goal is to train a general-purpose robot which can tackle multiple tasks with different object compositions in diverse environments. However, most current policy learning approaches with imitation learning or reinforcement learning can only learn to solve one task at a time, and usually operate on a fixed set of objects. To achieve human-level generalization, many efforts have been made on performing robotic manipulation tasks specified by natural language [34, 35, 56]. Language can not only bring its compositionality to low-level robot skills, but also operate as a high-level planner for long-horizon tasks.

*Work done during internship at UCSD.

†Work done outside of Amazon.

Recently, there is a trend on leveraging the pre-trained vision-language foundation models [38,47] as the backbone encoders for generalizing robot skills. For example, CLIPORT [56] uses the pre-trained CLIP model [47] as the image observation and language instruction encoder, and learn a manipulation policy on top for tasks across different object arrangements. While encouraging results have been shown in this line of research on generalization across training tasks, it is still very challenging for the learned policy to generalize to unseen tasks and environments. For example, as shown in Figure 1 (a), our experiments show that if we train such a policy on two tasks including (i) pack objects of different shapes in the brown box and (ii) put blocks of different colors in the bowls of different colors, it is very challenging for the policy to generalize to a task that put objects of different shapes in different colored bowls. Furthermore, the difficulty drastically increases when we need to perform this task in the real world with a different robot.

Our key insight is that, while the action generator (i.e., the policy) cannot generalize well, the action classifier with foundation models can still achieve a high accuracy even when “zero-shot” transferred to unseen environments [35, 56]. In this paper, we leverage vision-language foundation models to provide feedback during deploying the policy in unseen tasks and environments. We utilize the feedback from foundation models to fine-tune the policy following test-time training [28,57,61], which updates model parameters during test-time. Specifically, we propose Policy Adaptation from Foundation model Feedback (PAFF) with a play and relabel pipeline. When adapting a trained policy to a new task or new environment, we first let the policy **play**, that is, the model continuously generates and performs actions given a series of language instructions in the new task and we record the demonstrations including the visual observations and model’s actions. Of course, the instructions and the outcome demonstrations will often not match under the out-of-distribution environment. We then let the model **relabel** to make the correction, that is, the recorded demonstrations can be automatically relabeled by the vision-language pre-trained model. By taking the visual observations of recorded demonstrations as inputs, the pre-trained model can retrieve accurate language instructions correspondingly. Given the accurate paired demonstrations and instructions in the new environment, we can fine-tune and adapt the policy with them. We emphasize that the whole process of PAFF performs in an automatic way using trained models without human interventions.

We carefully design a broad range of language conditioned robotic adaptation experiments to evaluate the policy adaptation across object composition, tasks and environments including from simulation to the real world. Our evaluations consist of (i) **Compositional Generalization** in Fig. 1 (a), where we train a policy to pack objects of dif-

ferent shapes in the brown box, and put blocks of different colors in the bowls of different colors, and adapt it to put objects of different shapes in the bowls of different colors. (ii) **Out-of-distribution Generalization** in Fig. 1 (c), where we train a policy to pack certain objects in the brown box, and adapt it to unseen objects; and in Fig. 1 (d), where we adapt a policy trained on seen environments to an unseen environment with different textures and placements of static elements such as the sliding door, the drawer and the switch. (iii) **Sim-to-real Transfer** in Fig. 1 (b), where we adapt a policy trained on simulation data to the real-world. We show PAFF improves baselines by a large-margin in all evaluations. In sim-to-real transfer, our method significantly improves the success rate by an average of 49.6% on four tasks than the baseline. Our pipeline fills the domain gap between simulation and real world through utilizing the generalization capability of the foundation model. Our method also increases the success rate from 17.8% to 35.0% in the compositional generalization evaluation, and from 48.4% to 63.8% for packing unseen objects. When adapting the policy to an unseen environment, our method increases the success rate of completing 5 chains of language instructions from 5% to 11% over the baseline method. The extensive evaluation results show that PAFF can effectively adapt a language conditioned policy to unseen objects, tasks, environments, and realize sim-to-real transfer.

2. Related Work

Language Conditioned Manipulation. Instruction-based manipulation has been a popular research topic in robotics [4, 6, 7, 20, 33, 35, 37, 45, 54, 55] not only because it provides a user-friendly interface, but also because the compositional properties of language allows skill generalization and guides long-horizon planning. Recently, the advancement of foundation models [5, 9, 22, 36, 47, 49, 50, 58, 64] have led to significant progresses in generalizable manipulation skill learning [2, 25, 34, 56]. These approaches adopt the pre-trained foundation models to encode the language instruction and visual observation, and train policy network on top. While this pipeline successfully generalizes the policy across different tasks and new objects, it still has a hard time on generalizing to completely unseen environments and tasks. In this paper, we propose policy adaptation from foundation model feedback to adapt our policy during deployment in unseen environments and tasks.

Policy Adaptation with Visual Inputs. The ability to adapt a visuo-motor control policy to unseen environments is the key in many applications such as sim2real transfer. A popular way to achieve such generalization is using domain randomization [46, 51, 59] and data augmentation [8, 19, 29, 30, 69] for learning invariant visual representations. However, it is still very challenging for these policies to generalize beyond the randomization range of train-

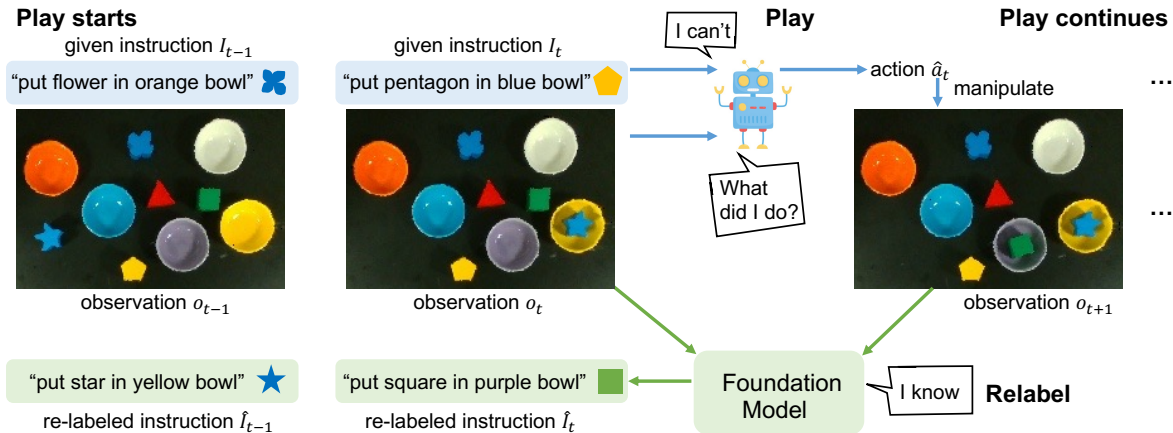


Figure 2. The pipeline of policy adaptation from foundation model feedback (PAFF). When we adapt a trained policy to a new task, we first let the robot **play**, that is, the policy continuously predicts and performs actions given a series of randomly generated language instructions. We record these demonstrations including visual observations and model’s actions. After that, we let the model **relabel**, that is, the vision-language foundation model relabels the demonstrations by retrieving the language instructions given the recorded visual observations. We then fine-tune the policy with the paired observations and instructions, and corresponding actions, which are collected in an automatic way.

ing. Another line of research is to utilize GANs [15, 73] to translate input images across domains before forwarding to the policy [24, 52, 70, 71]. However, all the approaches mentioned above require the tasks and environments are aligned across training and deployment. To adapt to a different task, researchers have proposed different ways to finetune the policy during deploying to the new environment [18, 26, 28, 39, 67, 72]. For example, Julian et al. [26] propose to collect 800 grasping attempts when transferring to a new robotic manipulation environment, and use the collected data to fine-tune the policy. However, this process requires significant human interventions. Hansen et al. [18] follows test-time training [57, 61] to allow the policy to continue training after deployment in a new environment by exploring self-supervision. Instead of using self-supervision, we utilize the feedback from vision-language foundation models to fine-tuning the policy in an automatic way.

Pre-trained Foundation Models in Robotics. Besides language-conditioned manipulation tasks, we have witnessed the utilization of foundation models on various robotic learning algorithms [1, 17, 40, 43, 48, 53, 65, 66, 68]. For example, Radosavovic et al. [48] propose to train a masked autoencoder [21] on large scale data such as Ego4D [16], and perform real robot learning with behavior cloning using this pre-trained representation. It is shown that such pre-training helps largely reduce the sample efficiency and improves the performance. However, it is not thoroughly studied how pre-training can affect out-of-distribution generalization, and most works only use the pre-trained network as policy encoder. In this work, besides utilizing the representations from foundation models, we also exploit their generalization ability in recognizing visual concepts when transferred to unseen environments to provide feedback during deployment in a new environment.

3. Method

We propose policy adaptation from foundation model feedback (PAFF) with a play and relabel pipeline. As summarized in Algorithm 1, our pipeline consists of two stages. In the first stage, we train a policy and fine-tune a vision-language pre-trained model with the training demonstrations. The second stage involves play and relabel in the new task. Specifically, as shown in Fig. 2, we first let the policy **play** with a series of randomly generated instructions and record the demonstrations including visual observations and actions. While the manipulation could be wrong, we let the fine-tuned vision-language model **relabel** the demonstrations by retrieving the language instructions given the visual observations. With the re-labeled pairs of demonstration-instruction data, we fine-tune the policy to adapt to the new task.

Our method utilizes the generalization capability of the vision-language foundation model in recognizing visual concepts, to annotate the data for fine-tuning the policy of the new task in an automatic way. PAFF can effectively adapt a language conditioned policy to manipulate unseen objects and solve new tasks in novel environments, and realize sim-to-real transfer.

3.1. Language Conditioned Policy

Inspired by previous work [34, 56], we consider the problem of learning a language conditioned policy π that outputs actions a_t given input $\gamma_t = (o_t, I_t)$ consisting of a visual observation o_t and a language instruction I_t as below:

$$\pi(\gamma_t) = \pi(o_t, I_t) \rightarrow a_t \quad (1)$$

We use an imitation-learning based method to learn the language conditioned policy. We experiment with two types

Algorithm 1 PAFF

Function: train a policy and fine-tune a foundation model
▷ **Stage 1**, Sec. 3.1 and Sec. 3.3

Input: training demonstrations; a randomly initialized policy π_θ ; a pre-trained foundation model f_ϕ

- 1: **for** each training demonstration **do**
- 2: optimize the policy with $\mathcal{L}(\pi_\theta(o_t, l_t), a_t)$
- 3: optimize the pre-trained foundation model with $\mathcal{L}(f_\phi(o_t, o_{t+1}, l_t))$
- 4: **end for**

Function: play and relabel in a new task, then fine-tune the policy
▷ **Stage 2**, Sec. 3.2

Input: initial observation; a series of language instructions

- 5: **for** each language instruction **do**
 - 6: predict action $\hat{a}_t = \pi_\theta(o_t, l_t)$, perform manipulation and record demonstrations
 - 7: **end for**
 - 8: **for** each recorded demonstration **do**
 - 9: retrieve a language instruction $\hat{l}_t = f_\phi(o_t, o_{t+1})$
 - 10: **end for**
 - 11: **for** each recorded demonstration **do**
 - 12: optimize the policy with $\mathcal{L}(\pi_\theta(o_t, \hat{l}_t), \hat{a}_t)$
 - 13: **end for**
-

of manipulation platforms with different action space: (i) We follow CLIPORT [56] to formulate tabletop object manipulation (*e.g.*, pick up a block and place it in a bowl) as a series of pick-and-place affordance predictions, where the objective is to detect actions and each action involves a start and final end-effector pose, and build our model upon a two-stream architecture in CLIPORT; (ii) We follow CALVIN [35] to formulate manipulation tasks that require continuous control (*e.g.*, push the sliding door to the left side) as 7-DoF control, and build our model based on a hierarchical architecture in HULC [34]. We adopt the same imitation training loss $\mathcal{L}(\pi_\theta(o_t, l_t), a_t)$ as defined in CLIPORT and HULC to optimize the policy.

3.2. Play and Relabel

After we train a policy and adapt it to a new task, it can often make mistakes (*e.g.* the policy picks up a star and places it in a yellow bowl given the instruction “put the flower in the orange bowl” as shown in Fig. 2). If our model can correct the instruction to a matching one (*e.g.* “put the star in the yellow bowl”), then the paired demonstration-instruction data can be used to fine-tune the policy.

We propose policy adaptation from foundation model feedback to collect the data for fine-tuning the policy of a new task in an automatic way, without human interventions. Specifically, as shown in Fig. 2, we first let the policy “play” with a series of randomly generated language instructions.




Figure 3. Given the language instruction in an unseen environment, the policy performs actions that do not match the instruction. The vision-language foundation model relabels the demonstration by taking the sequential visual observations as inputs and retrieving the language instruction among all possible instructions.

Given the current visual observation o_t and an instruction I_t , the policy predicts the action a_t and the robot performs the corresponding manipulation, reaching a new scene with the visual observation o_{t+1} . After a certain number of language instructions are executed, the scene will be reset automatically by the robot, which moves the objects out of the containers to the table following the instruction “move the objects out”. In this way, the policy can “play” in the new task, that is, the robot will explore the scene by continuously receiving language instructions and manipulating objects. We record these demonstrations including the visual observations $\{o_t\}_{t=1}^T$ and the robot’s actions $\{\hat{a}_t\}_{t=1}^T$ predicted by the trained policy.

After recording the demonstrations, we let the model relabel the demonstrations using the fine-tuned vision-language foundation model as shown in Fig. 2. We formulate the task of labeling the demonstrations as visual-to-language retrieval. Specifically, given the visual observations o_t and o_{t+1} , the foundation model retrieves a language instruction \hat{I}_t among all possible language instructions. Benefit from pre-training on large-scale data, the foundation model can generalize well across domains, thus is able to retrieve accurate language instructions for the recorded demonstrations.

After play and relabel, we collect new pairs of demonstration-instruction data automatically including the visual observations $\{o_t\}_{t=1}^T$, the re-labeled language instructions $\{\hat{I}_t\}_{t=1}^T$, and the actions $\{\hat{a}_t\}_{t=1}^T$, which are used to fine-tune the policy for the adaptation to the new task.

3.3. Vision-language Foundation Model

To make the model relabel the recorded demonstrations with the accurate instruction automatically, we cast the task of labeling demonstrations as visual-to-language re-

trieval, which is a mainstream downstream task in vision-language pre-training. The existing vision-language foundation model such as CLIP [47] pre-trains on single images and their captions describing the static content, thus can not be directly applied to label the recorded demonstrations with sequential visual observations. For example, for manipulations that require continuous control such as “pushing the sliding door to the left side” as shown in Fig. 3, in order to recognize this action, the model requires the capability of both spatial and temporal reasoning among the sequential visual observations. The demonstration can not be labeled if the model lacks the ability to infer temporal information of the visual observations.

We utilize a Spatio-Temporal Adapter (ST-Adapter) [42] to fine-tune CLIP for the ability to reason about sequential visual observations of the recorded demonstrations. Specifically, the pre-trained CLIP f_ϕ adopts a visual encoder f_v for visual representations, and a text encoder f_t for language representations. We add a depth-wise 3D convolution layer [12] between the transformer [10] layers of the CLIP visual encoder f_v , so that f_v can take sequential visual observations o_t and o_{t+1} as inputs to conduct spatio-temporal modeling (We denote sequential observations of executing a language instruction as o_t and o_{t+1} for simplicity). We use the CLIP text encoder f_t to extract representations of the language instruction l_t . We adopt contrastive learning to maximize the similarity between the representations of visual observations denoted as $f_v(o_t, o_{t+1})$ and representations of the corresponding language instruction denoted as $f_t(l_t)$, and minimize the similarity between $f_v(o_t, o_{t+1})$ and representations of other language instructions with Noise-Contrastive Estimation (NCE) [41],

$$\mathcal{L}(f_\phi(o_t, o_{t+1}, l_t)) = \text{NCE}(f_v(o_t, o_{t+1}), f_t(l_t)) \quad (2)$$

$$\text{NCE}(x_i, y_i) = -\log \frac{\exp(x_i^T y_i / \tau)}{\sum_{j=1}^N \exp(x_i^T y_j / \tau)} \quad (3)$$

where N is the number of the language instructions and τ is the temperature hyper-parameter and we set $\tau = 0.05$. We train the 3D convolution layers on the the seen tasks and freeze the parameters of the pre-trained CLIP.

After fine-tuning the vision-language foundation model, given a recorded demonstration, we calculate the similarity between the representations of the visual observations and the representation of all language instructions, and retrieve the language instruction with the maximum similarity. The fine-tuned foundation model can achieve a high retrieval accuracy even when “zero-shot” transferred to unseen environments. For example, our fine-tuned CLIP achieves a retrieval accuracy of 99.3% on the demonstrations of the unseen environment in the CALVIN [35] benchmark. Through utilizing the generalization capability of the large-scale vision-language foundation model in recognizing visual concepts, our method can relabel the recorded demonstrations for fine-tuning the policy in an automatic way.

4. Experiment

4.1. Evaluation Settings

We carefully design four language conditioned policy adaptation evaluations based on CLIPORT [56] and CALVIN [35], to evaluate (i) **Compositional Generalization**, where we train a policy to pack objects of different shapes in the brown box (“pack-shapes”), and put blocks of different colors in the bowls of different colors (“put-blocks-in-bowls”), and adapt it to put objects of different shapes in the bowls of different colors (“put-shapes-in-bowls”) as shown in Fig. 1 (a); (ii) **Out-of-distribution Generalization**, where we train a policy on packing seen objects and adapt it to unseen objects (“pack-unseen-objects”) using Google Scanned Objects dataset [11] as shown in Fig. 1 (c), with the same split as CLIPORT, and train a policy on seen environments and adapt it to a new environment with different textures and differently positioned static elements such as the sliding door and light button in CALVIN as shown in Fig. 1 (d). We use environment A, B and C for training, and environment D for adaptation; (iii) **Sim-to-real Transfer**, where we train a policy on simulation data and adapt it to real world with four tasks including “pack-blocks”, “packing-shapes”, “put-blocks-in-bowls”, “put-shapes-in-bowls” as shown in Fig. 1 (b). We also explore a challenging task (*i.e.* compositional “put-shapes-in-bowls”), where the policy is trained on simulation data of “packing-shapes” and “put-blocks-in-bowls”.

For the compositional and out-of-distribution evaluations in the CLIPORT platform, we follow CLIPORT to report task success rate of 100 evaluation instances on 10 different scenes (with different blocks, objects and bowls), where the success rate is the number of the correctly placed objects, divided by the total number of the objects. For the out-of-distribution evaluation in the CALVIN platform, we follow CALVIN to evaluate Long-Horizon Multi-Task Language Control (LH-MTLC), which treats the 34 tasks as subtasks and evaluates 100 unique instruction chains, each consisting of five sequential tasks. The policy receives the next subtask in a chain only if it successfully completes the current one. We calculate the success rate of each task in the chain and the averaged successful sequence length as the evaluation metrics. For the sim-to-real transfer evaluation, we report task success rate of 10 evaluation instances for each task, where a instance consists of executing 5 language instructions.

4.2. Implementation Details

To train policies, we follow CLIPORT [56] and HULC [34], but adopt the pre-trained MDETR [27] as the visual and language encoder (See Sec. 4.5 for details). To fine-tune the vision-language foundation model (*i.e.*, CLIP [47]) for labeling, in the CLIPORT platform, we use

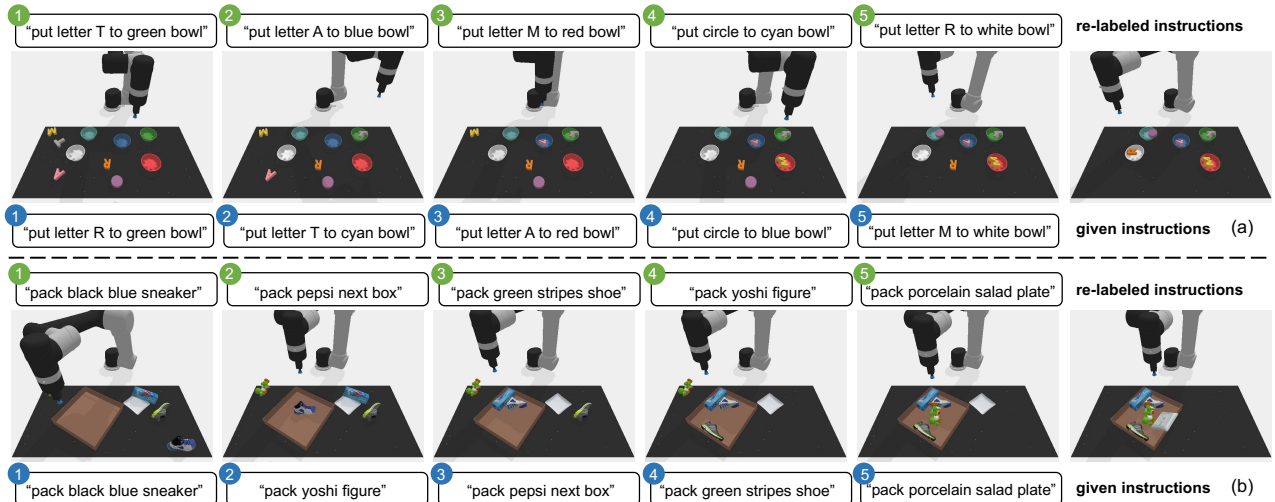


Figure 4. Given instructions vs. automatically re-labeled instructions of the recorded demonstrations with five consecutive steps, where we let the policy play in new tasks. The policy performs manipulations, where the outcome demonstrations do not match the instructions. By taking the recorded demonstrations as inputs, our model retrieves accurate language instructions correspondingly.

Table 1. Results of compositional and out-of-distribution generalization evaluation in the CLIPORT platform. The evaluation metric is the success rate, where each step receives a new instruction in the left column, while a step receives a new instruction when the previous instruction is executed correctly in the right column.

Method	put-shapes-in-bowls		pack-unseen-objects	
CLIPORT [56]	28.0%	16.8%	58.9%	46.1%
MdetrORT [27]	33.8%	17.8%	62.0%	48.4%
AugORT [44]	34.4%	18.9%	63.1%	49.0%
Ours	51.0%	35.0%	72.8%	63.8%

the start and end frame of the visual observations as the input of CLIP. In the CALVIN platform, a language instruction in the training set corresponds to sequential observations with 64 frames, and we sample 8 frames following TSN [62] as the input of CLIP. We freeze the pre-trained CLIP and train ST-Adapter [42] for 300 epochs with the training data of seen tasks and environments on 8 GPUs.

For evaluations in the CLIPORT platform, we first train a policy with 100 demonstrations per task for 200 epochs on a single GPU. We then let the robot play by following language instructions generated by a template “put _ in _”. We record 40 demonstrations continuously for each scene, where a demonstration performs five instructions. After the model relabels demonstrations, we save recorded demonstrations, where the similarity score between representations of visual observations and retrieved instruction is higher than 3.0. We fine-tune the policy on the saved demonstrations for 100 epochs.

For the evaluation in the CALVIN platform, we train a policy on the training data of environments A, B, C for 10 epochs on 16 GPUs. We then let the robot play by following instructions that are randomly chosen from all instructions used in seen environments. We record 500 demonstra-

Table 2. Results of out-of-distribution generalization evaluation in the CALVIN platform. The number denotes the success rate of each subtask in the chain, and “Len” denotes the averaged successful sequence length.

Method	1	2	3	4	5	Len
HULC [34]	43%	14%	4%	1%	0%	0.62
MdetrLC [27]	69%	38%	20%	7%	4%	1.38
AugLC [44]	69%	43%	22%	9%	5%	1.48
Ours	72%	47%	30%	13%	11%	1.73

tions in the environment D, where a demonstration performs five subtasks. We fine-tune the policy on the re-labeled and saved demonstrations for 5 epochs.

4.3. Baselines

Besides CLIPORT [56] and HULC [34], we also adopt two methods MdetrORT and MdetrLC as the baselines by replacing the visual and language encoder in CLIPORT and HULC with the pre-trained MDETR [27]. They train a policy of seen tasks in exactly the same way as our method without play and relabel. Based on MdetrORT and MdetrLC, we further adopt two methods AugORT and AugLC, which use data augmentation to train a policy following Pashkevich et al. [44].

4.4. Results

4.4.1 Compositional Generalization

As shown in Fig. 4, this kind of adaptation across object composition is challenging. For example, given the language instruction “put letter T to cyan bowl” in the second column, the trained policy wrongly picks up letter A object and puts it in the blue bowl in the third column. Such a mistake can be corrected by our fine-tuned vision-language

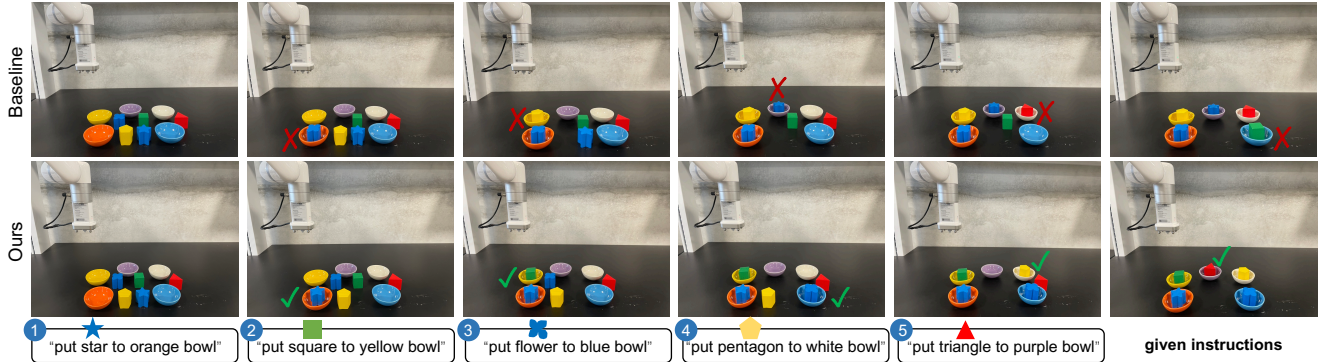


Figure 5. Real world experiments for sim-to-real transfer with five consecutive steps. The task puts objects of the specified shape in the bowls of the specified color.

Table 3. Results of the sim-to-real transfer evaluation. The evaluation metric is the success rate, where each step receives a new language instruction. “compositional” denotes that the policy is trained on “pack-shapes” and “put-blocks-in-bowls”.

Method	pack-blocks	pack-shapes	put-blocks-in-bowls	put-shapes-in-bowls	put-shapes-in-bowls (compositional)
MdetrORT [27]	72%	32%	40%	24%	14%
Ours	98%	92%	88%	82%	70%

foundation model as it retrieves “put letter A to blue bowl” as the language instruction given the visual observations of the second and the third column.

We show evaluation results of compositional generalization in Tab. 2. Our method achieves the best results under both evaluation protocols. Compared with MdetrORT [27], which trains a policy on seen tasks in exactly the same way as ours, our method improves the success rate by 17.2%, which shows the effectiveness of our play and re-label pipeline for compositional generalization. The fine-tuned foundation model can be generalized to recognize the visual concepts (*i.e.* objects and bowls) in the compositional setting as shown in Fig. 4 (a), thus the re-labeled demonstrations can be used to fine-tune the policy in the new task. We further observe that MdetrORT performs better than CLIPORT by replacing the pre-trained CLIP [47] with the pre-trained MDETR [27] as the visual and language encoder. MDETR learns object-aware representations by aligning features of visual regions and text phrases, and can benefit the learning of a object manipulation policy. AugORT further improves performance slightly by applying data augmentation [44] during training, but still achieves poor results since compositional generalization is outside the randomization range of training.

4.4.2 Out-of-distribution Generalization

We design two experiments to evaluate the out-of-distribution generalization. Tab. 2 lists the results of packing unseen objects in the CLIPORT platform. Our method surpasses the baseline methods under both evaluation protocols. As shown in Fig. 4, when the policy picks up an object

that does not match the language instruction, the vision-language foundation model can correct the mistake by retrieving a corresponding language instruction given the visual observations. Since the foundation model is pre-trained on large scale data, it can be generalized to recognizing unseen objects in the new task.

Tab. 1 reports the results of adapting a policy to an unseen environment in the CALVIN platform. Our method outperforms the baseline methods in terms of the success rate of each subtask in the chain and the averaged successful sequence length. As shown in Fig. 3, given the sequential visual observations, our fine-tuned vision-language foundation model is able to relabel the recorded demonstrations through spatial and temporal reasoning with the Spatio-Temporal Adapter [42]. We further observe that using the pre-trained MDETR as the visual and language encoder in HULC also improves performance, and applying the data augmentation helps adapt the policy to an unseen environment. But their results still lag far behind ours.

4.4.3 Sim-to-real Transfer

Adapting a policy trained on the simulation data to the real world is challenging because of the huge simulation-to-real gap such as different textures, lighting, colors, and objects. We evaluate four tasks in Tab. 3, and our method significantly boosts performance in each task compared with the baseline method (without play and relabel). We further evaluate a more difficult setting, which also involves the evaluation of compositional generalization by training policy on the simulation data of “pack-shapes” and “put-blocks-in-bowls”. As shown in Fig. 5, when we directly

Table 4. Ablation Study on the number of the recorded demonstrations. Each demonstration consists of executing five language instructions. We evaluate the compositional generalization and the out-of-distribution generalization in the CLIPORT platform.

Number	put-shapes-in-bowls		pack-unseen-objects	
10	36.0%	21.3%	69.7%	59.0%
20	37.6%	24.6%	73.2%	62.4%
40	51.0%	35.0%	72.8%	63.8%

adapt the learned policy to the real world, it picks up the incorrect object and often places it in the bowl with wrong color. By contrast, through play and relabel, our method can perform manipulation correctly given the language instructions. Our pipeline fills the domain gap between simulation and real world through utilizing the generalization capability of the pre-trained foundation model.

4.5. Ablation Study

Number of recorded demonstrations. We study the effect of the number of the recorded demonstrations when we let the policy “play” in the new task, where the policy executes five language instructions in each demonstration and then follows the instruction “move the objects out” to reset the scene automatically. As shown in Tab. 4, in general, recording 40 demonstrations for each scene achieves the best performance for both the compositional generalization and the out-of-distribution generalization evaluation. Recording more demonstrations increases the diversity of the collected data for fine-tuning the policy, thus benefits the policy adaptation to the new task.

Visual and language encoder We explore different pre-trained foundation models including MoCo [22], DenseCL [63], MAE [21], GLIP [31], CLIP [47] and MDETR [27] as the visual and language encoder to train a policy, and evaluate their performances without play and relabel on the compositional generalization experiment with two evaluation protocols. For image pre-trained model, we use the text encoder in CLIP as the language encoder. The results are listed in Tab. 5. We have some observations. First of all, Vit [10] based MAE achieves the worst result. Compared with other models that use ResNet [23] and Swin Transformer [32] for multi-scale feature maps, Vit produces feature maps at a single scale, which does not benefit object manipulation policy learning. Second, DenseCL surpasses MoCo through performing dense pair-wise contrastive learning at the level of pixels rather than contrast global features. DenseCL tailors the self-supervised learning method for dense prediction tasks such as object detection, which is beneficial to learning a policy for object manipulation. Furthermore, GLIP and MDETR outperforms CLIP by aligning features of regions in the image and phrases in the text. The learned object-level and language-aware visual representations contributes to a better manipu-

Table 5. Ablation study on the visual and language encoder using different foundation models. We evaluate the compositional generalization with “put-shapes-in-bowls” in the CLIPORT platform.

MoCo [22]	DenseCL [63]	MAE [21]	GLIP [31]	CLIP [47]	MDETR [27]
22.6%	26.4%	19.0%	30.1%	28.0%	33.8%
13.2%	13.6%	11.8%	17.3%	16.8%	17.8%

Table 6. Ablation study on the temporal reasoning mechanism for fine-tuning the foundation model. We evaluate the out-of-distribution generalization in the CALVIN platform with the averaged successful sequence length as the evaluation metric.

2D joint- attention [3]	3D joint- attention [13, 60]	Divided- attention [14]	ST-Adapter [42]
1.60	1.53	1.46	1.73

lation policy. Finally, MDETR achieves better results than GLIP, and we adopt this pre-trained foundation model as our visual and language encoder.

Temporal reasoning mechanism. To capitalize the foundation model CLIP pre-trained on single images and captions for labeling recorded demonstrations with sequential visual observations, we explore different temporal reasoning mechanism to fine-tune CLIP including (i) 2d joint-attention [3], which uses 2d convolution to flatten each image and performs joint-attention [10], (ii) 3d joint-attention [13, 60], which uses 3d convolution to flatten two images and performs joint-attention, (iii) divided-attention [14], which adds temporal attention [3] among different images, and (iv) Spatio-Temporal Adapter (ST-Adapter) [42], which adds a depth-wise 3D convolution layer [12] between each transformer layer. As listed in Tab. 6, ST-Adapter achieves the best result on the out-of-distribution generalization evaluation in the CALVIN platform. ST-Adapter trains the depth-wise 3D convolution layers for spatio-temporal reasoning, and meanwhile freezes the parameters of the pre-trained CLIP to reserve the generalization capability, thus can retrieve more accurate language instructions to fine-tune the policy.

5. Conclusion

In this work, we propose policy adaptation from foundation model feedback (PAFF), which leverages the vision-language foundation model to collect the data for fine-tuning the policy in unseen tasks and environments automatically. We evaluate our method on a broad range of language conditioned policy adaptation experiments including compositional generalization, out-of-distribution generalization and sim-to-real transfer, and show great superiority of our method.

Acknowledgements. This project was supported, in part, by Amazon Research Award and gifts from Qualcomm.

References

- [1] Rishabh Agarwal, Max Schwarzer, Pablo Samuel Castro, Aaron Courville, and Marc G Bellemare. Reincarnating reinforcement learning: Reusing prior computation to accelerate progress. In *Advances in Neural Information Processing Systems*, 2022. 3
- [2] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Chuyuan Fu, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, Daniel Ho, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Eric Jang, Rosario Jauregui Ruano, Kyle Jeffrey, Sally Jesmonth, Nikhil Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Kuang-Huei Lee, Sergey Levine, Yao Lu, Linda Luu, Carolina Parada, Peter Pastor, Jornell Quiambao, Kanishka Rao, Jarek Rettinghouse, Diego Reyes, Pierre Sermanet, Nicolas Sievers, Clayton Tan, Alexander Toshev, Vincent Vanhoucke, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Mengyuan Yan, and Andy Zeng. Do as i can and not as i say: Grounding language in robotic affordances. In *arXiv preprint arXiv:2204.01691*, 2022. 2
- [3] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, volume 2, page 4, 2021. 8
- [4] V Blukis, RA Knepper, and Y Artzi. Few-shot object grounding for mapping natural language instructions to robot control. In *CoRL*, 2020. 2
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 2
- [6] Yiye Chen, Ruinian Xu, Yunzhi Lin, and Patricio A Vela. A joint network for grasp detection conditioned on natural language commands. In *ICRA*, pages 4576–4582. IEEE, 2021. 2
- [7] Yiye Chen, Ruinian Xu, Yunzhi Lin, and Patricio A Vela. A joint network for grasp detection conditioned on natural language commands. In *ICRA*, pages 4576–4582. IEEE, 2021. 2
- [8] Karl Cobbe, Oleg Klimov, Chris Hesse, Taehoon Kim, and John Schulman. Quantifying generalization in reinforcement learning, 2018. 2
- [9] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 5, 8
- [11] Laura Downs, Anthony Francis, Nate Koenig, Brandon Kinman, Ryan Hickman, Krista Reymann, Thomas B McHugh, and Vincent Vanhoucke. Google scanned objects: A high-quality dataset of 3d scanned household items. *arXiv preprint arXiv:2204.11918*, 2022. 5
- [12] Christoph Feichtenhofer. X3d: Expanding architectures for efficient video recognition. In *CVPR*, pages 203–213, 2020. 5, 8
- [13] Christoph Feichtenhofer, Haoqi Fan, Yanghao Li, and Kaiming He. Masked autoencoders as spatiotemporal learners. *arXiv preprint arXiv:2205.09113*, 2022. 8
- [14] Yuying Ge, Yixiao Ge, Xihui Liu, Dian Li, Ying Shan, Xiaohu Qie, and Ping Luo. Bridging video-text retrieval with multiple choice questions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16167–16176, 2022. 8
- [15] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 3
- [16] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18995–19012, 2022. 3
- [17] Agrim Gupta, Stephen Tian, Yunzhi Zhang, Jiajun Wu, Roberto Martín-Martín, and Li Fei-Fei. Maskvit: Masked visual pre-training for video prediction. *arXiv preprint arXiv:2206.11894*, 2022. 3
- [18] Nicklas Hansen, Rishabh Jangir, Yu Sun, Guillem Alenyà, Pieter Abbeel, Alexei A. Efros, Lerrel Pinto, and Xiaolong Wang. Self-supervised policy adaptation during deployment. In *International Conference on Learning Representations*, 2021. 3
- [19] Nicklas Hansen and Xiaolong Wang. Generalization in reinforcement learning by soft data augmentation. In *International Conference on Robotics and Automation*, 2021. 2
- [20] Jun Hatori, Yuta Kikuchi, Sosuke Kobayashi, Kuniyuki Takahashi, Yuta Tsuboi, Yuya Unno, Wilson Ko, and Jethro Tan. Interactively picking real-world objects with unconstrained spoken language instructions. In *ICRA*, pages 3774–3781. IEEE, 2018. 2
- [21] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. 3, 8
- [22] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9729–9738, 2020. 2, 8
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. 8
- [24] Daniel Ho, Kanishka Rao, Zhuo Xu, Eric Jang, Mohi Khansari, and Yunfei Bai. Retinagan: An object-aware approach to sim-to-real transfer. In *ICRA*, pages 10920–10926. IEEE, 2021. 3
- [25] Eric Jang, Alex Irpan, Mohi Khansari, Daniel Kappler, Fredrik Ebert, Corey Lynch, Sergey Levine, and Chelsea Finn.

- Bc-z: Zero-shot task generalization with robotic imitation learning. In *CoRL*, pages 991–1002. PMLR, 2022. [2](#)
- [26] Ryan Julian, Benjamin Swanson, Gaurav S Sukhatme, Sergey Levine, Chelsea Finn, and Karol Hausman. Efficient adaptation for end-to-end vision-based robotic manipulation. 2020. [3](#)
- [27] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetr-modulated detection for end-to-end multi-modal understanding. In *ICCV*, pages 1780–1790, 2021. [5](#), [6](#), [7](#), [8](#)
- [28] Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. Rma: Rapid motor adaptation for legged robots. *arXiv preprint arXiv:2107.04034*, 2021. [2](#), [3](#)
- [29] Michael Laskin, Kimin Lee, Adam Stooke, Lerrel Pinto, Pieter Abbeel, and Aravind Srinivas. Reinforcement learning with augmented data. *arXiv preprint arXiv:2004.14990*, 2020. [2](#)
- [30] Kimin Lee, Kibok Lee, Jinwoo Shin, and Honglak Lee. A simple randomization technique for generalization in deep reinforcement learning. *ArXiv*, abs/1910.05396, 2019. [2](#)
- [31] Liunian Harold Li, Pengchuan Zhang, Haotian Zhang, Jianwei Yang, Chunyuan Li, Yiwu Zhong, Lijuan Wang, Lu Yuan, Lei Zhang, Jenq-Neng Hwang, et al. Grounded language-image pre-training. In *CVPR*, pages 10965–10975, 2022. [8](#)
- [32] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021. [8](#)
- [33] Corey Lynch and Pierre Sermanet. Language conditioned imitation learning over unstructured data. *arXiv preprint arXiv:2005.07648*, 2020. [2](#)
- [34] Oier Mees, Lukas Hermann, and Wolfram Burgard. What matters in language conditioned robotic imitation learning. *arXiv preprint arXiv:2204.06252*, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)
- [35] Oier Mees, Lukas Hermann, Erick Rosete-Beas, and Wolfram Burgard. Calvin: A benchmark for language-conditioned policy learning for long-horizon robot manipulation tasks. *IRAL*, 2022. [1](#), [2](#), [4](#), [5](#)
- [36] Antoine Miech, Jean-Baptiste Alayrac, Lucas Smaira, Ivan Laptev, Josef Sivic, and Andrew Zisserman. End-to-end learning of visual representations from uncurated instructional videos. In *CVPR*, pages 9879–9889, 2020. [2](#)
- [37] Dipendra K Misra, Jaeyong Sung, Kevin Lee, and Ashutosh Saxena. Tell me dave: Context-sensitive grounding of natural language to manipulation instructions. *IJRR*, 35(1-3):281–300, 2016. [2](#)
- [38] Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. Slip: Self-supervision meets language-image pre-training. In *ECCV*, pages 529–544. Springer, 2022. [2](#)
- [39] Anusha Nagabandi, Ignasi Clavera, Simin Liu, Ronald S Fearing, Pieter Abbeel, Sergey Levine, and Chelsea Finn. Learning to adapt in dynamic, real-world environments through meta-reinforcement learning. *arXiv preprint arXiv:1803.11347*, 2018. [3](#)
- [40] Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhinav Gupta. R3m: A universal visual representation for robot manipulation. *arXiv preprint arXiv:2203.12601*, 2022. [3](#)
- [41] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. [5](#)
- [42] Junting Pan, Ziyi Lin, Xiatian Zhu, Jing Shao, and Hongsheng Li. St-adapter: Parameter-efficient image-to-video transfer learning for action recognition. *arXiv preprint arXiv:2206.13559*, 2022. [5](#), [6](#), [7](#), [8](#)
- [43] Simone Parisi, Aravind Rajeswaran, Senthil Purushwalkam, and Abhinav Gupta. The unsurprising effectiveness of pre-trained vision models for control. *arXiv preprint arXiv:2203.03580*, 2022. [3](#)
- [44] Alexander Pashevich, Robin Strudel, Igor Kalevtykh, Ivan Laptev, and Cordelia Schmid. Learning to augment synthetic images for sim2real policy transfer. In *IROS*, pages 2651–2657. IEEE, 2019. [6](#), [7](#)
- [45] Chris Paxton, Yonatan Bisk, Jesse Thomason, Arunkumar Byravan, and Dieter Fox. Prospecion: Interpretable plans from language by predicting the future. In *ICRA*, pages 6942–6948. IEEE, 2019. [2](#)
- [46] Lerrel Pinto, Marcin Andrychowicz, Peter Welinder, Wojciech Zaremba, and Pieter Abbeel. Asymmetric actor critic for image-based robot learning. *arXiv preprint arXiv:1710.06542*, 2017. [2](#)
- [47] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763. PMLR, 2021. [2](#), [5](#), [7](#), [8](#)
- [48] Ilija Radosavovic, Tete Xiao, Stephen James, Pieter Abbeel, Jitendra Malik, and Trevor Darrell. Real-world robot learning with masked visual pre-training. *arXiv preprint arXiv:2210.03109*, 2022. [3](#)
- [49] Jack W Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John Aslanides, Sarah Henderson, Roman Ring, Susannah Young, et al. Scaling language models: Methods, analysis & insights from training gopher. *arXiv preprint arXiv:2112.11446*, 2021. [2](#)
- [50] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67, 2020. [2](#)
- [51] Fabio Ramos, Rafael Possas, and Dieter Fox. Bayessim: Adaptive domain randomization via probabilistic inference for robotics simulators. *Robotics: Science and Systems XV*, Jun 2019. [2](#)
- [52] Kanishka Rao, Chris Harris, Alex Irpan, Sergey Levine, Julian Ibarz, and Mohi Khansari. RL-cycleGAN: Reinforcement learning aware simulation-to-real. In *CVPR*, pages 11157–11166, 2020. [3](#)
- [53] Younggyo Seo, Kimin Lee, Stephen L James, and Pieter Abbeel. Reinforcement learning with action-free pre-training from videos. In *International Conference on Machine Learning*, pages 19561–19579. PMLR, 2022. [3](#)

- [54] Mohit Shridhar and David Hsu. Interactive visual grounding of referring expressions for human-robot interaction. *arXiv preprint arXiv:1806.03831*, 2018. 2
- [55] Mohit Shridhar and David Hsu. Interactive visual grounding of referring expressions for human-robot interaction. *arXiv preprint arXiv:1806.03831*, 2018. 2
- [56] Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Cliport: What and where pathways for robotic manipulation. In *CoRL*, pages 894–906. PMLR, 2022. 1, 2, 3, 4, 5, 6
- [57] Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *International conference on machine learning*, pages 9229–9248. PMLR, 2020. 2, 3
- [58] Romal Thoppilan, Daniel De Freitas, Jamie Hall, Noam Shazeer, Apoorv Kulshreshtha, Heng-Tze Cheng, Alicia Jin, Taylor Bos, Leslie Baker, Yu Du, et al. Lamda: Language models for dialog applications. *arXiv preprint arXiv:2201.08239*, 2022. 2
- [59] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep 2017. 2
- [60] Zhan Tong, Yibing Song, Jue Wang, and Limin Wang. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *arXiv preprint arXiv:2203.12602*, 2022. 8
- [61] Dequan Wang, Evan Shelhamer, Shaoteng Liu, Bruno Olshausen, and Trevor Darrell. Tent: Fully test-time adaptation by entropy minimization. In *International Conference on Learning Representations*, 2020. 2, 3
- [62] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks for action recognition in videos. *IEEE transactions on pattern analysis and machine intelligence*, 41(11):2740–2755, 2018. 6
- [63] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *CVPR*, pages 3024–3033, 2021. 8
- [64] Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*, 2021. 2
- [65] Ted Xiao, Harris Chan, Pierre Sermanet, Ayzaan Wahid, Anthony Brohan, Karol Hausman, Sergey Levine, and Jonathan Tompson. Robotic skill acquisition via instruction augmentation with vision-language models. In *Workshop on Language and Robotics at CoRL*, 2022. 3
- [66] Tete Xiao, Ilija Radosavovic, Trevor Darrell, and Jitendra Malik. Masked visual pre-training for motor control. *arXiv preprint arXiv:2203.06173*, 2022. 3
- [67] Yifan Xu, Nicklas Hansen, Zirui Wang, Yung-Chieh Chan, Hao Su, and Zhuowen Tu. On the feasibility of cross-task transfer with model-based reinforcement learning. *arXiv preprint arXiv:2210.10763*, 2022. 3
- [68] Yanjie Ze, Nicklas Hansen, Yinbo Chen, Mohit Jain, and Xiaolong Wang. Visual reinforcement learning with self-supervised 3d representations. *arXiv preprint arXiv:2210.07241*, 2022. 3
- [69] Amy Zhang, Rowan McAllister, Roberto Calandra, Yarin Gal, and Sergey Levine. Learning invariant representations for reinforcement learning without reconstruction. *arXiv preprint arXiv:2006.10742*, 2020. 2
- [70] Jingwei Zhang, Lei Tai, Peng Yun, Yufeng Xiong, Ming Liu, Joschka Boedecker, and Wolfram Burgard. Vr-goggles for robots: Real-to-sim domain adaptation for visual control. *IEEE Robotics and Automation Letters*, 4(2):1148–1155, 2019. 3
- [71] Qiang Zhang, Tete Xiao, Alexei A Efros, Lerrel Pinto, and Xiaolong Wang. Learning cross-domain correspondence for control with dynamics cycle-consistency. *arXiv preprint arXiv:2012.09811*, 2020. 3
- [72] Wenxuan Zhou, Lerrel Pinto, and Abhinav Gupta. Environment probing interaction policies. *arXiv preprint arXiv:1907.11740*, 2019. 3
- [73] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*, pages 2223–2232, 2017. 3