
648 A TRAINING APPROACH AND PSEUDOCODE
649

650 We utilize Proximal Policy Optimization (PPO) (Schulman et al., 2017) to train our policy network.
651 PPO is an on-policy, actor-critic deep RL algorithm. The optimization objective for the policy is as
652 follows:
653

654
655
$$\mathcal{L}_{\text{policy}}(\theta) = \hat{\mathbb{E}} \left[\min \left(r_t(\theta) \hat{A}_t, \text{clip} (r_t(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_t \right) \right] \quad (9)$$

656

657 Here, \hat{A}_t denotes the estimation of the advantage function, and $r_t(\theta)$ represents the probability ratio,
658 defined as $r_t(\theta) = \frac{\pi_\theta(a_t | s_t)}{\pi_{\theta_{\text{old}}}(a_t | s_t)}$, where π_θ represents the new policy and $\pi_{\theta_{\text{old}}}$ represents the old policy.
659

660 In optimizing the conditional policy module, we utilize the $\mathcal{L}_{\text{policy}}$ loss function. Notably, similar
661 to the approach utilized in VariBAD (Zintgraf et al., 2019), the optimization of the task inference
662 module does not rely on the $\mathcal{L}_{\text{policy}}$ loss function. Instead, we adopt a composite loss function that
663 combines reconstruction and contrastive learning objectives. The specific pseudo-code is shown in
664 Algorithm 1.
665

666
667 **Algorithm 1** CFOHLR
668

669 **Require:** Encoder q_ϕ and Decoder p_θ of VAE; Coarse policy π_θ ; Fine policy π_ω ; Weight generator
670 W_α Skill-specific expert models $\text{MoE}_{\theta_i}^{i=0, \dots, K}$; Hypernetworks H_ϕ ; VAE buffer \mathcal{D}_{VAE} ; Policy
671 buffer $\mathcal{D}_{\text{Policy}}$; The number of meta-episodes n_{meta} ; The number of rollout episodes per meta-
672 episode n_{roll} ; Language instruction \mathcal{U} .
673 **while** $i = 0, \dots, N_{\text{update}}$ **do**
674 Sample K training tasks $M_i^{i=0, \dots, K} \sim M_{\text{train}}$
675 **for** timestep $t = 0, \dots, n_{\text{roll}} * H - 1$ **do**
676 **if** $t \bmod H = 0$ **then**
677 Reset rollout episode for each task, obtain $S_t = \{s_{t,1}, s_{t,2}, \dots, s_{t,n}\}$
678 **end if**
679 **for** $j = 0, \dots, K$ **do**
680 Obtain weights $\alpha_{1,j}, \dots, \alpha_{k,j} = W_\alpha(\mathcal{U}_j)$ for each skill-specific expert module.
681 Obtain the output of the coarse policy π_θ , denoted as $O_{\text{MOE}} = \sum_{i=0}^K \alpha_i \cdot \text{MoE}_{\theta_i}$.
682 Leverage H_ϕ to generate the network parameters of the fine policy, $\pi_\omega = H_\phi(z_t^j)$.
683 Obtain the action $a_{t,j} = \pi_\omega(O_{\text{MOE}})$.
684 **end for**
685 Finally, obtain actions for each task $A_t = \{a_{t,1}, a_{t,2}, \dots, a_{t,n}\}$.
686 Take an environment step, obtaining $S_{t+1} = \{s_{t+1,1}, s_{t+1,2}, \dots, s_{t+1,n}\}$ and $R_t = \{r_{t+1,1}, r_{t+1,2}, \dots, r_{t+1,n}\}$.
687 Add the transition $(S_t, A_t, R_{t+1}, S_{t+1})$ to \mathcal{D}_{VAE} and $\mathcal{D}_{\text{Policy}}$.
688 Update task representations $Z_{t+1} = \{z_{t+1,n} = q_\phi(\tau_{:t+1,n})\}_{i=0, \dots, n}$.
689 **end for**
690 Update VAE by minimizing $\mathcal{L} = \mathcal{L}_{\text{VAE}} + \mathcal{L}_{\text{contra}}$
691 Update policy θ, ω and weight generator α by minimizing $L_{\text{actor}} + L_{\text{critic}}$.
692 **end while**
693
694

695 B LIMITATIONS AND FUTURE WORK
696

697 Despite the significant progress, our method has limitations that were not addressed in this study.
698 Notably, it is not directly applicable to the cross-entity adaptation problem, which involves general-
699 izing a policy from one robotic entity to another. This limitation affects the overall generalizability
700 of the policy. Future research will focus on tackling the challenge of cross-entity adaptation in a
701 zero-shot manner, thereby enhancing the policy generalization.
702

702 **C IMPLEMENTATION DETAILS**

703

704 **C.1 REFERENCE IMPLEMENTATIONS**

705

706 **SDVT, LDM, and VariBAD** We adapt the SDVT (Lee et al., 2023), LDM (Lee & Chung, 2021),
 707 and VariBAD (Zintgraf et al., 2019) algorithms to the Meta-World benchmark. These algorithms
 708 are all based on the VariBAD method, which itself is grounded in the Bayesian Adaptive MDP
 709 (BAMDP) framework. VariBAD employs a VAE architecture consisting of a recurrent encoder and
 710 a dynamics decoder to obtain task representations. LDM introduces a virtual training procedure to
 711 VariBAD to address out-of-distribution challenges. Building on LDM, SDVT uses a Gaussian mix-
 712 ture distribution to model the latent space of the VAE. Notably, the virtual training steps of the LDM
 713 and SDVT methods are included in the total count of training steps, as these virtual processes neces-
 714 sitate agent interaction with the environment to obtain real states for generating imagined samples.
 715 We used open-source code to reproduce the results of the SDVT, LDM, and VariBAD methods, re-
 716 spectively, available at <https://github.com/suyoung-lee/SDVT>, <https://github.com/suyoung-lee/LDM>, and <https://github.com/lmzintgraf/varibad>.

717

718 **Million** Million (Bing et al., 2023) introduces a meta-RL paradigm comprising an instruction
 719 phase and a trial phase, integrating transformers with language instruction to improve task adap-
 720 tation capabilities. We used open-source code to reproduce the results of the Million methods,
 721 respectively, available at <https://github.com/yaoxt3/MILLION>.

722 **C.2 HYPERPARAMETERS**

723

724 **C.2.1 SDVT**

725

726 We used the default hyperparameters from the paper, which are shown in Table 3.

727

728 **C.2.2 LDM AND VARIBAD**

729

730 We used the default hyperparameters from the paper, which are shown in Table 4.

731 **C.2.3 MILLION**

732

733 We used the default hyperparameters from the paper, which are shown in Table 5.

734

735 **C.2.4 OURS**

736

737 **C.3 NETWORK ARCHITECTURE**

738

739 Our method utilizes a context-based architecture, comprising a task inference module and a condi-
 740 tional policy module. For the task inference module, similar to SDVT, we also employ a Gaussian
 741 mixture VAE to model the latent space. This module consists of an RNN-based encoder and a
 742 prediction decoder. Before being input into the encoder or decoder, all state, action, and reward
 743 inputs pass through embedding networks. Regarding the conditional policy module, it includes
 744 language-selected, skill-specific expert modules and a hypernetwork-based, task-aware policy. Sim-
 745 ilarly, before being input into the conditional policy module, all state, action, and reward inputs pass
 746 through embedding networks.

747 **C.4 TASK DESCRIPTIONS**

748

749 In Table 11, we provide the language instructions for each of the 50 Meta-World tasks.

750

751 **D DETAILED EXPERIMENTAL RESULTS**

752

753 We adhere to the success criterion established by Meta-World. A timestep is considered successful
 754 when the distance between the task-relevant object and the target falls within an acceptable range.
 755 Furthermore, an entire rollout episode is deemed successful if the agent achieves success at any
 756 timestep during the episode.

Table 3: Hyperparameters used for Garage experiments with SDVT

Description	ML10	ML45
Meta-Task Hyperparameters		
Meta-batch size	10	10
Tasks sampled per epoch	10	10
General Hyperparameters		
Batch size	1,000	1,000
Path length per roll-out	1,000	1,000
Discount factor	0.99	0.99
Algorithm-Specific Hyperparameters		
Policy hidden sizes	(256, 256)	(256, 256)
Activation function	tanh	tanh
Policy learning rate	7×10^{-4}	7×10^{-4}
PPO epochs num	5	5
VAE learning rate	1×10^{-3}	1×10^{-3}
Latent dimension	5	5
PPO num minibatches	10	10
PPO clip param	0.1	0.1
Policy num steps	5	5
Size of VAE buffer	1,000	1,000
KL weight	0.1	0.1
VAE mixture num	10	10
Gaussian loss coefficient	1.0	1.0
Action embedding size	16	16
State embedding size	32	32
Reward embedding size	16	16
Virtual ratio increment	0.05	0.05
Number of virtual skills	3	3
RL loss through encoder	False	False
VAE loss coefficient	1.0	1.0

D.1 PERFORMANCE ON INDIVIDUAL TASKS

D.1.1 ML-10

D.1.2 ML-45

D.2 LEARNING CURVES

In Figure 5, we present the mean and standard deviation of returns and success rates across five random seeds.

E ADDITIONAL RESULTS

E.1 VISULIZATIONS

To demonstrate the quality of the learned task representations, we employed t-SNE Van der Maaten & Hinton (2008) to map the task representation vectors into a 2D space, enabling the visualization of these representations. For each testing task, 150 transitions from the meta-testing phase were sampled to visualize the task representations. As depicted in Figure 6, our method effectively distinguishes task representations from different categories, with additional separation observed among tasks within the same category.

810
 811
 812
 813 Table 4: The hyperparameters used in experiments with LDM and VariBAD are consistent across
 814 both models in the general and policy categories of SDVT, as outlined in Table 3. The only difference
 815 lies in the modeling of the latent space: SDVT utilizes a Gaussian mixture model, while both LDM
 816 and VariBAD employ a Gaussian model.

Description	ML10	ML45
Meta-Task Hyperparameters		
Meta-batch size	10	10
Tasks sampled per epoch	10	10
General Hyperparameters		
Batch size	1,000	1,000
Path length per roll-out	1,000	1,000
Discount factor	0.99	0.99
Algorithm-Specific Hyperparameters		
VAE learning rate	1×10^{-3}	1×10^{-3}
Latent dimension	5	5
Size of VAE buffer	1,000	1,000
KL weight	0.1	0.1
Gaussian loss coefficient	1.0	1.0
VAE loss coefficient	1.0	1.0

827
 828
 829
 830
 831
 832
 833
 834
 835
 836
 837
 838
 839
 840 Table 5: Hyperparameters used in experiments with Million.
 841

Description	ML10	ML45
Meta-Task Hyperparameters		
Meta-batch size	10	10
Tasks sampled per epoch	10	10
General Hyperparameters		
Batch Timesteps	1,000	1,000
Action repeat	1,000	1,000
Demonstration action	1,000	1,000
repeat		
Max trials per episode	750	750
Discount factor	0.99	0.99
Algorithm-Specific Hyperparameters		
Learning rate	$1e - 4$	$1e - 4$
GAE lambda	0.97	0.97
Epsilon eta	1×10^{-2}	1×10^{-2}
Epsilon alpha	1×10^{-2}	1×10^{-2}
Epsilon alpha mu	0.0075	0.0075
Epsilon alpha sigma	$1e - 5$	$1e - 5$

864

865

866

Table 6: Hyperparameters used in experiments with Ours.

867

Description	ML10	ML45
Meta-Task Hyperparameters		
Meta-batch size	10	10
Tasks sampled per epoch	10	10
General Hyperparameters		
Batch size	1,000	1,000
Path length per roll-out	1000	
Discount factor	0.99	
Algorithm-Specific Hyperparameters		
Policy hidden sizes	(256, 256)	(256, 256)
Activation function	tanh	tanh
Policy learning rate	7×10^{-4}	7×10^{-4}
PPO epochs num	5	5
VAE learning rate	1×10^{-3}	1×10^{-3}
Latent dimension	5	5
PPO num minibatches	10	10
PPO clip param	0.1	0.1
Policy num steps	5	5
RL loss through encoder	False	False
Action embedding size	16	16
State embedding size	32	32
Reward embedding size	16	16
Size of VAE buffer	1,000	1,000
KL weight	0.1	0.1
VAE mixture num	10	10
Gaussian loss coefficient	1.0	1.0
VAE loss coefficient	1.0	1.0
Decode reward	True	True
Decode state	True	True
Weight of holistic contrastive	0.01	0.01
Weight of local contrastive	0.01	0.01

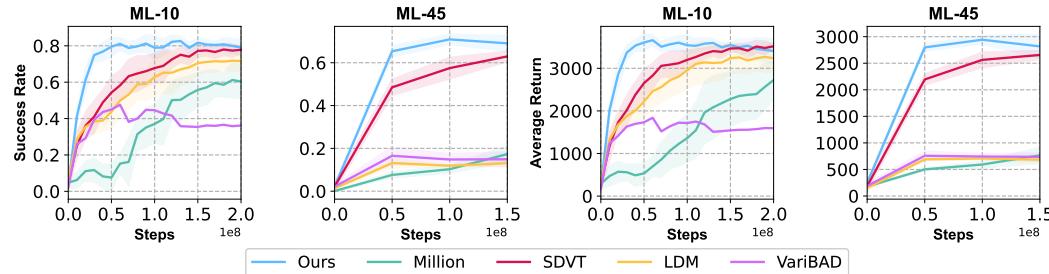
900

901

902

903

904



912

913

914

Figure 5: **Learning Curves on ML-10 and ML-45.** The maximum success rates and corresponding returns of our methods, along with baseline comparisons, are presented. The plot shows the mean and standard deviation of returns across five random seeds.

915

916

917

Table 7: ML-10 of Meta-World success rate (%). We present the final success rates of our method and baseline approaches on both the training and test tasks of the Meta-World ML-10 benchmark. All results are reported as the mean success rate $\pm 95\%$ confidence interval of five seeds.

Index. Task	Ours	w/o C2F	w/o HLC	SDVT	Million	LDM	VariBAD
1. Reach	85.2±3.6	53.2±4.6	44.0±7.3	53.6±14.4	10.4±11.3	50.4±7.9	78.0±5.7
2. Push	86.0±3.0	70.4±9.8	74.8±3.0	74.0±3.8	44.5±17.8	31.2±11.2	2.4±2.0
3. Pick-place	85.2±4.5	66.0±6.9	66.0±2.9	53.2±6.0	48.1±29.8	37.2±20.0	0.8±0.7
4. Door-open	81.2±1.9	99.6±0.6	97.2±3.9	100.0±0.0	81.1±25.1	99.6±0.6	74.8±25.4
5. Drawer-close	86.8±1.4	100.0±0.0	100.0±0.0	100.0±0.0	56.1±32.1	100.0±0.0	100.0±0.0
6. Button-press	84.4±3.1	100.0±0.0	100.0±0.0	99.6±0.6	80.0±27.7	98.4±1.0	88.4±4.1
7. Peg-insert-side	88.0±2.5	48.8±18.6	62.0±13.1	52.0±5.0	21.5±18.3	26.4±15.3	0.0±0.0
8. Window-open	86.4±4.4	99.6±0.6	100.0±0.0	100.0±0.0	80.0±27.7	99.6±0.6	96.8±1.1
9. Sweep	89.6±2.7	93.2±2.1	93.2±3.5	89.2±3.8	77.6±24.0	92.4±3.3	0.0±0.0
10. Basketball	82.8±4.5	93.6±4.0	92.0±5.6	72.8±15.5	36.1±31.0	89.2±4.0	0.0±0.0
Train mean	85.6±3.9	82.4±7.7	82.9±4.8	79.4±7.1	53.5±40.0	72.4±11.4	44.1±7.9
11. Drawer-open	86.4±3.4	78.0±24.1	80.0±12.6	72.8±28.5	0.0±0.0	92.8±8.1	51.6±21.1
12. Door-close	87.2±3.8	74.4±25.7	81.6±15.7	76.4±19.7	97.5±2.8	26.0±37.7	71.6±22.4
13. Shelf-place	82.4±5.5	1.6±3.1	0.4±0.8	0.0±0.0	0.3±0.5	0.4±0.8	0.0±0.0
14. Sweep-into	90.0±5.1	67.6±34.0	82.8±8.8	64.8±12.2	17.5±5.0	57.6±31.9	4.8±3.2
15. Lever-pull	82.4±3.4	5.2±7.4	0.4±0.8	3.2±3.2	13.7±26.9	0.4±0.8	0.4±0.8
Test mean	85.7±4.9	45.4±13.7	49.0±5.5	43.4±9.4	25.8±9.1	35.4±17.1	25.7±7.5

Table 8: ML-10 of Meta-World returns. We present the performance metrics of our method and baseline approaches on both the training and test tasks of the Meta-World ML-10 benchmark. All results are reported as the mean return $\pm 95\%$ confidence interval of five seeds.

Index. Task	Ours	w/o C2F	w/o HLC	SDVT	Million	LDM	VariBAD
1. Reach	3704±149	3778±128	3520±141	3763±296	2324±447	3668±285	4054±138
2. Push	3769±135	3338±342	4094±90	3675±272	2225±750	1795±812	63±28
3. Pick-place	3742±127	2089±254	2420±116	1712±125	1678±803	1258±709	7±1
4. Door-open	3740±91	4503±51	4313±78	4439±26	2790±980	4442±47	2978±470
5. Drawer-close	3708±75	4857±7	4811±30	4852±10	2505±1302	4809±67	4637±77
6. Button-press	3622±144	3489±93	3250±108	3372±60	1337±734	3226±60	2028±155
7. Peg-insert-side	3703±113	2359±678	2827±421	2443±266	1179±697	1364±773	9±1
8. Window-open	3787±157	4479±49	4398±49	4476±51	2331±942	4384±61	3692±202
9. Sweep	3786±147	4093±85	3963±100	3801±208	2849±932	3997±189	92±28
10. Basketball	3705±185	3532±196	3576±202	2937±618	1624±774	3433±196	9±2
Train mean	3727±221	3652±289	3717±112	3547±203	2084±1365	3238±613	1757±178
11. Drawer-open	3796±112	2477±393	2477±190	2660±396	1876±384	2697±475	2036±329
12. Door-close	3740±119	2489±566	2887±769	3087±944	3302±415	1272±1538	2113±558
13. Shelf-place	3746±157	492±246	607±115	341±99	141±204	309±272	0±0
14. Sweep-into	3751±128	1619±786	1705±434	1444±564	716±264	1200±793	172±96
15. Lever-pull	3634±197	305±44	251±29	285±60	208±44	278±59	324±41
Test mean	3734±165	1476±224	1585±290	1563±418	1249±205	1151±692	929±208

972
973
974 Table 9: ML-45 of Meta-World success rate (%). We present the final success rates of our method
975 and baseline approaches on both the training and test tasks of the Meta-World ML-45 benchmark.
976 All results are reported as the mean success rate $\pm 95\%$ confidence interval of five seeds.
977

Index.	Task	Ours	w/o C2F	w/o HLC	SDVT	Million	LDM	VariBAD
1.	Assembly	68.0±2.0	0.5±0.3	0.5±0.3	0.5±0.3	0.0±0.0	0.0±0.0	0.0±0.0
2.	Basketball	67.5±1.3	22.0±5.7	21.5±1.7	38.0±5.2	0.0±0.0	0.0±0.0	0.0±0.0
3.	Button-press-topdown	65.0±3.2	99.0±0.6	98.5±0.3	98.0±0.7	0.0±0.0	9.5±2.8	27.5±6.7
4.	Button-press-topdown-wall	72.0±2.3	96.0±2.3	98.5±0.6	99.0±0.3	0.0±0.0	9.5±2.7	17.0±4.6
5.	Button-press-wall	75.5±3.0	94.0±2.1	95.0±1.8	100.0±0.0	43.2±10.8	30.0±3.9	27.0±6.1
6.	Button-press-wall	71.0±2.4	88.0±3.1	76.0±4.8	80.5±4.2	49.2±6.0	41.0±9.7	32.0±4.8
7.	Coffee-button	65.0±2.8	100.0±0.0	100.0±0.0	98.5±0.9	42.3±14.7	55.5±11.0	33.5±8.7
8.	Coffee-pull	73.5±2.9	62.5±1.3	70.0±3.6	37.0±5.8	0.2±0.1	1.0±0.3	0.0±0.0
9.	Coffee-push	76.5±1.5	57.5±10.7	69.5±5.5	37.5±4.9	30.8±9.7	11.5±3.3	7.5±2.1
10.	Dial-turn	73.5±2.6	73.0±6.1	75.0±4.9	77.5±2.6	62.2±2.4	8.5±1.7	26.0±6.4
11.	Disassemble	74.0±3.2	69.5±13.6	79.0±5.3	78.0±4.3	0.0±0.0	0.5±0.3	4.5±2.6
12.	Door-close	72.5±1.3	100.0±0.0	99.5±0.3	100.0±0.0	51.0±11.9	98.5±0.9	65.0±13.8
13.	Door-open	74.5±2.8	92.0±3.3	94.0±2.7	95.5±0.9	0.0±0.0	0.0±0.0	0.0±0.0
14.	Drawer-close	70.5±1.7	96.0±1.4	100.0±0.0	99.5±0.3	100.0±0.0	99.0±0.6	100.0±0.0
15.	Drawer-open	68.0±2.4	99.5±0.3	95.0±1.5	98.5±0.6	0.0±0.0	15.0±6.1	10.0±2.4
16.	Faucet-open	71.0±2.7	98.0±0.5	96.0±2.0	98.5±0.9	11.2±3.8	30.5±1.5	54.0±9.7
17.	Faucet-close	77.0±1.0	98.5±0.9	87.5±3.8	99.5±0.3	42.0±13.5	22.0±4.4	28.5±7.1
18.	Hammer	75.5±1.5	8.5±5.0	33.0±11.1	0.0±0.0	12.2±7.1	1.5±0.9	2.0±0.8
19.	Handle-press-side	69.0±1.4	100.0±0.0	99.0±0.6	100.0±0.0	14.2±4.9	6.0±2.2	38.5±8.5
20.	Handle-press	70.0±1.3	100.0±0.0	99.5±0.3	100.0±0.0	65.5±3.2	45.5±2.9	57.0±2.5
21.	Handle-pull-side	72.5±3.0	96.0±2.0	98.0±1.2	89.5±3.6	17.7±4.6	0.0±0.0	1.5±0.6
22.	Handle-pull	70.5±1.9	76.5±13.3	99.5±0.3	63.5±12.2	27.0±10.1	11.5±0.3	1.0±0.6
23.	Lever-pull	76.5±1.2	55.5±9.6	9.5±5.6	52.0±10.4	0.0±0.0	0.0±0.0	1.5±0.3
24.	Peg-insert-side	72.0±1.7	9.0±1.1	33.5±2.5	3.5±0.6	0.0±0.0	0.0±0.0	0.0±0.0
25.	Pick-place-wall	72.0±1.1	61.5±4.2	59.5±6.4	46.0±3.5	13.5±5.8	0.0±0.0	0.5±0.3
26.	Pick-out-of-hole	67.0±3.6	39.0±9.9	52.0±5.0	53.5±5.7	0.0±0.0	0.0±0.0	0.0±0.0
27.	Push	79.5±3.0	38.0±3.0	48.0±2.2	27.5±3.2	7.7±2.9	42.0±3.8	46.0±5.8
28.	Push-back	66.5±1.9	69.0±5.7	79.5±3.1	64.0±4.2	0.0±0.0	0.0±0.0	1.0±0.6
29.	Push	74.5±2.6	43.0±4.2	68.0±4.9	38.5±7.5	8.3±2.4	4.5±1.0	3.0±1.0
30.	Pick-place	67.5±1.2	50.5±2.4	58.5±2.8	47.5±3.6	10.3±4.3	0.0±0.0	1.0±0.3
31.	Plate-slide-side	69.0±2.8	47.5±4.8	48.0±3.7	67.0±3.8	36.7±7.0	0.0±0.0	0.5±0.3
32.	Plate-slide-side	71.0±3.2	95.0±1.0	91.5±2.2	92.5±2.3	0.0±0.0	0.5±0.3	13.5±7.9
33.	Plate-slide back	67.5±1.5	96.5±0.6	98.5±0.9	91.5±1.5	0.0±0.0	1.5±0.6	4.5±1.5
34.	Plate-slide-back-side	75.5±2.4	80.5±5.1	89.0±2.5	83.5±2.7	0.0±0.0	0.5±0.3	6.0±2.0
35.	Peg-unplug-side	69.5±1.7	66.5±3.7	76.0±4.7	53.5±3.8	5.3±1.3	6.5±2.8	4.5±1.5
36.	Soccer	73.0±1.7	21.5±4.8	18.0±1.2	26.0±2.4	10.5±2.1	9.0±2.2	8.0±1.3
37.	Stick-push	72.0±1.4	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0
38.	Stick-pull	73.5±1.9	0.0±0.0	0.5±0.3	0.0±0.0	0.0±0.0	0.0±0.0	0.0±0.0
39.	Push-wall	72.0±2.4	54.5±6.4	88.0±1.1	54.0±10.0	7.8±3.1	0.5±0.3	1.0±0.6
40.	Reach-wall	75.0±0.6	40.5±6.9	60.0±4.1	26.5±6.1	13.3±5.3	49.5±5.6	75.0±5.2
41.	Shelf-place	74.5±2.5	6.0±2.8	1.0±0.6	1.0±0.6	0.8±0.5	0.0±0.0	0.0±0.0
42.	Sweep-into	66.0±1.4	94.5±1.2	95.5±1.3	81.5±8.1	25.0±4.7	9.0±1.9	11.0±2.6
43.	Sweep	73.0±1.1	74.5±2.9	83.0±3.6	36.5±12.3	0.0±0.0	0.0±0.0	0.0±0.0
44.	Window-open	64.5±0.7	99.0±0.3	98.5±0.9	100.0±0.0	60.0±8.7	7.9	13.5±3.6
45.	Window-close	67.0±2.7	100.0±0.0	99.5±0.3	99.5±0.3	11.8±4.0	17.0±3.2	29.5±8.0
Train mean		71.4±4.3	66.0±5.8	69.8±4.1	63.0±5.5	17.3±8.9	14.2±2.6	17.2±4.2
46.	Bin-picking	76.0±3.2	1.0±1.0	1.5±0.9	3.5±2.6	0.2±0.3	0.0±0.0	0.0±0.0
47.	Box-close	70.0±4.3	1.0±1.0	6.5±3.9	0.5±0.9	1.7±2.9	0.5±0.9	0.5±0.9
48.	Hand-insert	69.5±2.6	0.5±0.9	2.5±2.6	3.5±3.6	7.5±7.7	3.0±3.4	3.0±2.3
49.	Door-lock	73.0±10.0	82.0±3.8	70.0±17.2	61.5±13.4	34.3±4.8	41.5±10.1	34.5±12.6
50.	Door-unlock	68.5±3.6	81.0±7.1	84.5±9.7	66.0±15.5	14.2±24.8	21.0±10.0	37.0±15.0
Test mean		71.4±3.5	33.1±3.0	33.0±6.3	27.0±8.9	11.6±7.1	13.2±6.0	15.0±6.5

1026

1027

1028

Table 10: ML-45 of Meta-World returns. We present the final returns of our method and baseline approaches on both the training and test tasks of the Meta-World ML-45 benchmark. All results are reported as the mean return $\pm 95\%$ confidence interval of five seeds.

1031

1032

1033

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043

1044

1045

1046

1047

1048

1049

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079

Index. Task	Ours	w/o C2F	w/o HLC	SDVT	Million	LDM	VariBAD
1. Assembly	2898±97	2847±27	2529±88	2590±61	329±53	154±27	101±13
2. Basketball	2885±101	1417±116	1444±107	1569±169	281±86	5±0	11±2
3. Button-press-topdown	2693±116	3582±100	3712±27	3586±50	884±128	988±104	1182±84
4. Button-press-topdown-wall	2950±69	3541±120	3686±52	3594±62	868±132	977±99	1209±81
5. Button-press-wall	3186±28	3143±111	3072±48	3193±101	553±102	667±68	615±96
6. Button-press-wall	3118±73	3344±108	3170±64	3275±44	495±108	711±156	633±87
7. Coffee-button	2746±53	3465±71	2731±336	3309±96	684±249	204±13	211±25
8. Coffee-pull	3123±79	1209±45	1385±128	877±118	58±4	40±2	40±4
9. Coffee-push	3197±39	1175±203	1463±193	729±55	228±62	41±4	65±16
10. Dial-turn	2861±70	3711±134	3396±108	3607±197	670±30	942±98	803±76
11. Disassemble	2990±85	2823±431	2811±291	2878±254	124±20	156±11	130±10
12. Door-close	2975±93	4310±44	4328±63	4481±19	1138±161	4359±27	2661±580
13. Door-open	3080±50	4004±116	3948±76	4010±109	636±61	624±76	607±62
14. Drawer-close	2936±35	4443±125	4730±28	4748±22	3625±462	4375±92	4482±122
15. Drawer-open	2855±112	4638±6	4065±99	4391±14	1746±95	1284±71	1356±127
16. Faucet-open	2945±94	4608±9	4276±158	4636±17	1669±47	2212±73	2584±299
17. Faucet-close	3074±23	4594±31	3997±173	4533±42	2349±214	2193±129	2174±185
18. Hammer	3013±68	516±28	1299±301	468±5	563±56	394±27	397±25
19. Handle-press-side	2949±27	4689±52	4707±29	4783±8	480±68	489±71	1377±270
20. Handle-press	2941±53	4648±66	4601±48	4579±64	2063±73	1791±100	2126±116
21. Handle-pull-side	2959±56	3647±143	3060±275	3340±240	165±64	19±2	24±2
22. Handle-pull	3000±83	3482±342	4019±59	2996±260	996±392	40±8	87±23
23. Lever-pull	2999±46	762±79	381±33	878±89	276±13	240±10	232±11
24. Peg-insert-side	2978±42	1307±81	1616±128	1238±50	192±36	11±0	10±1
25. Pick-place-wall	3102±60	2542±157	2728±136	1812±98	491±181	0±0	2±0
26. Pick-out-of-hole	2808±91	781±183	1206±142	922±122	23±5	10±1	13±1
27. Push	3157±49	2839±131	3313±108	2823±142	1555±125	3105±126	3193±89
28. Push-back	2897±111	1284±88	1799±66	881±203	16±2	7±1	5±0
29. Push	3094±115	2428±123	3421±56	2247±92	651±156	55±6	60±8
30. Pick-place	2921±54	1632±76	2165±83	1449±141	291±115	8±0	10±1
31. Plate-slide-side	2942±63	2516±138	2207±77	3221±98	1929±195	359±18	546±40
32. Plate-slide-side	3022±114	3250±84	2873±85	3784±178	818±54	191±31	662±139
33. Plate-slide back	2764±46	4235±28	4109±20	4165±52	1021±52	556±15	561±67
34. Plate-slide-back-side	3021±29	3776±179	4173±24	4124±58	631±142	183±35	728±134
35. Peg-unplug-side	2935±36	1390±217	1984±187	890±103	43±9	33±3	27±1
36. Soccer	2974±67	1056±25	1079±41	1052±65	515±152	278±30	321±16
37. Stick-push	3074±21	612±206	289±166	26±8	125±28	11±1	14±2
38. Stick-pull	2980±88	289±92	82±43	16±3	129±37	11±1	12±1
39. Push-wall	2925±74	2224±97	3586±50	2548±308	854±218	33±1	48±10
40. Reach-wall	3063±41	2747±263	3447±180	2330±273	1602±167	2906±263	3509±46
41. Shelf-place	3041±76	878±70	838±105	899±36	69±35	0±0	1±0
42. Sweep-into	2756±79	3962±140	4061±91	3095±422	879±157	238±43	207±45
43. Sweep	2902±59	2976±139	3251±110	1640±421	325±62	65±8	83±12
44. Window-open	2684±13	4142±51	4125±87	4225±31	710±62	471±30	795±116
45. Window-close	2828±112	4392±39	4397±28	4406±49	710±128	928±36	1152±96
Train mean	2961±149	2797±212	2879±201	2685±154	766±352	719±63	779±171
46. Bin-picking	2993±138	84±48	133±38	104±35	20±9	20±6	15±1
47. Box-close	2844±133	159±28	255±78	145±19	231±72	248±20	209±34
48. Hand-insert	2708±238	221±41	258±67	273±111	118±89	82±63	124±102
49. Door-lock	2960±335	2048±106	1901±747	1565±227	1659±184	1692±338	1589±374
50. Door-unlock	2962±51	1844±222	1952±206	1763±264	787±217	944±167	1219±339
Test mean	2912±105	871±117	900±226	770±140	563±60	597±154	631±203

1080

1081

1082

1083

1084

1085

1086

1087

Table 11: A list of all of the Meta-World tasks and a description of each task.

Task	Language instructions
assembly	pick up a nut and place it onto a peg
basketball	pick the basketball and place at the goal point
button-press-topdown	push the button down to the goal point
button-press-topdown-wall	bypass a wall and press a button from the top
button-press	press a button
button-press-wall	bypass a wall and press a button
coffee-button	push a button on the coffee machine
coffee-pull	place cup away
coffee-push	push cup to the goal point
dial-turn	rotate a dial 180 degrees
disassemble	pick a nut out of a peg
door-close	push the door to the goal point
door-open	pull the door to the goal point
drawer-close	push the drawer to the goal point
drawer-open	pull the drawer to the goal point
faucet-open	rotate the faucet counter-clockwise
faucet-close	rotate the faucet clockwise
hammer	push to the goal point with hammer
handle-press-side	press a handle down sideways
handle-press	press a handle down
handle-pull-side	pull a handle up sideways
handle-pull	pull a handle up
lever-pull	pull the lever to the goal point
peg-insert-side	insert the peg to the goal point
pick-place-wall	pick a puck, bypass a wall and place the puck
pick-out-of-hole	pick up a puck from a hole
reach	reach the goal point
push-back	push the puck back to the goal point
push	push the puck to the goal point
pick-place	pick the puck and place at the goal point
plate-slide	push the plate to the goal point
plate-slide-side	push the plate left to the goal point
plate-slide-back	push the plate back to the goal point
plate-slide-back-side	push the plate right to the goal point
peg-unplug-side	pull a peg sideways to the goal point
soccer	push a ball to the goal point
stick-push	grasp a stick and push a box using the stick
stick-pull	grasp a stick and pull a box with the stick
push-wall	bypass a wall and push a puck to a goal
reach-wall	bypass a wall and reach a goal
shelf-place	pick the puck and place on shelf at the goal point
sweep-into	sweep the puck into the box
sweep	sweep the puck off the table
window-open	push the window to the goal point
window-close	push the window to the goal point
bin-picking	grasp the puck from one bin and place it into another bin
box-close	grasp the cover and close the box with it
hand-insert	insert the gripper into a hole
door-lock	rotate the lock clockwise
door-unlock	rotate the lock counter-clockwise

1129

1130

1131

1132

1133

1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149

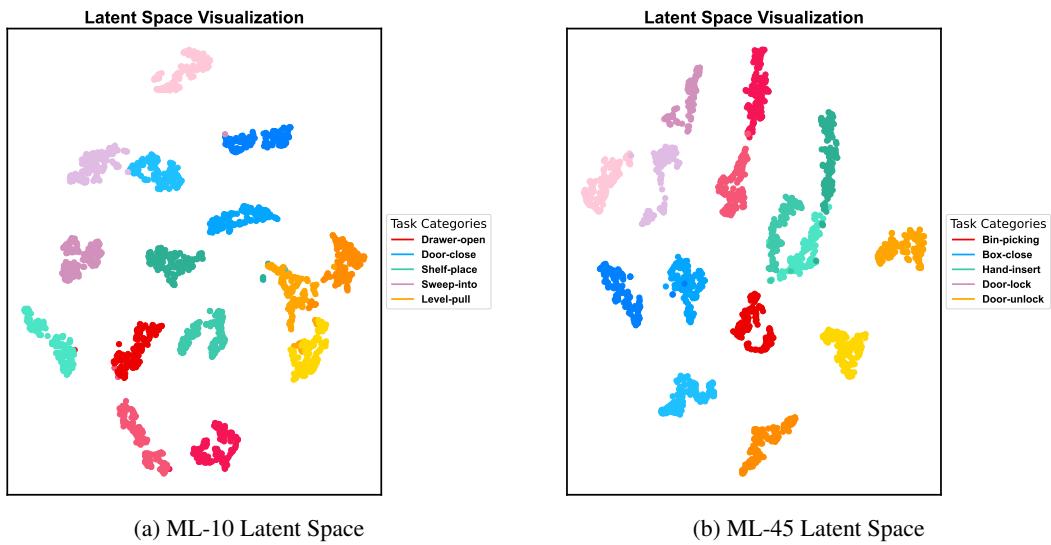


Figure 6: Latent Space Visualization. The t-SNE visualization of the learned task representation space for the ML-10 testing tasks is presented. We sampled three tasks from each task category of the test tasks, with each color scheme representing a different task category. Each point in the visualization corresponds to a task representation vector extracted from transitions and is color-coded according to the task properties.

1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187