

524 **A Pseudocode**

Algorithm 1 Subtask Decomposition and Virtual Training (SDVT)

Initialize Encoder q_ϕ , decoder p_θ , policy π_ψ , set of training tasks $\mathcal{M}_{\text{train}}$, virtual ratio p_v , GMVAE buffer \mathcal{B}_{VAE} , policy buffer \mathcal{B}_{pol} , total number of meta-episodes to train on n_{meta} , number of rollout episodes per meta-episode n_{roll} , running mean of subtask composition \bar{y} .

for meta-episode $k = 0, \dots, n_{\text{meta}} - 1$ **do**

- Sample a training task $M_k \sim p(\mathcal{M}_{\text{train}})$
- Sample a virtual meta-episode flag $V \sim \text{Bernoulli}(p_v)$
- if** $V = 1$ **then**

 - Sample an imaginary subtask composition $\tilde{y} \sim \text{Dirichlet}(\bar{y})$
 - end if**
 - Reset $h_0, y_0, \mathcal{B}_{\text{pol}}$
 - for** timestep $t = 0, \dots, n_{\text{roll}} \times H - 1$ **do**

 - if** $t \bmod H = 0$ **then**

 - Reset rollout episode $s_t \sim T_{0,k}(\cdot)$
 - end if**
 - Sample an action $a_t \sim \pi_\psi(\cdot | s_t, \mu_{\phi_z}(h_t, y_t), \sigma_{\phi_z}(h_t, y_t), \omega_{\phi_y}(h_t))$
 - Take an environment step

 - $s_{t+1} \sim T_k(\cdot | s_t, a_t)$
 - $r_{t+1} \leftarrow R_k(s_t, a_t, s_{t+1})$

 - if** $V = 1$ **then**

 - $\tilde{z}_t \sim \mathcal{N}(\mu_{\phi_z}(h_t, \tilde{y}), \sigma_{\phi_z}^2(h_t, \tilde{y}))$
 - $\tilde{h}_t = p_{\theta_h}(\tilde{z}_t)$
 - $\tilde{r}_{t+1} \sim p_{\theta_R}(\cdot | \text{DO}[s_t, a_t, s_{t+1}], \tilde{h}_t)$
 - Replace $r_{t+1} \leftarrow \tilde{r}_{t+1}$

 - else**

 - Add the transition $(s_t, a_t, s_{t+1}, r_{t+1})$ to \mathcal{B}_{VAE} \dots VAE NOT TRAINED WITH VIRTUAL DYNAMICS
 - end if**
 - Add the transition $(s_t, a_t, s_{t+1}, r_{t+1})$ to \mathcal{B}_{pol}
 - Update hidden embedding $h_{t+1} = q_{\phi_h}(\tau_{:t+1})$
 - Update subtask composition $y_{t+1} \sim \text{Cat}(\omega_{\phi_y}(h_{t+1}))$
 - Update the running mean of subtask composition $\bar{y} \leftarrow \text{RunningMeanUpdate}(\bar{y}, y_{t+1})$

 - end for**
 - Update GMVAE $\{\phi, \theta\} \leftarrow \{\phi, \theta\} + \nabla_{\{\phi, \theta\}} \sum_{t=0}^{H^+} \text{ELBO}_t$ with samples from \mathcal{B}_{VAE}
 - Update policy $\psi \leftarrow \psi + \nabla_\psi \mathcal{J}_{\text{pol}}^+$ with samples from \mathcal{B}_{pol}
 - Anneal virtual ratio $p_v \leftarrow p_v + \Delta p_v$

end for

525 **B ELBO Derivation**

526 The GMVAE's ELBO objective is derived as follows.

$$\begin{aligned}
 \mathbb{E}_{d(M_k, \tau_{:H^+})} [\log p_\theta(\tau_{:H^+})] &= \mathbb{E}_{d(M_k, \tau_{:H^+})} \left[\log \mathbb{E}_{q_\phi(y_t, z_t | h_t)} \left[\frac{p_\theta(\tau_{:H^+}, y_t, z_t)}{q_\phi(y_t, z_t | h_t)} \right] \right] \\
 &\geq \mathbb{E}_{d(M_k, \tau_{:H^+})} \left[\mathbb{E}_{q_\phi(y_t, z_t | h_t)} \left[\log \frac{p_\theta(\tau_{:H^+}, y_t, z_t)}{q_\phi(y_t, z_t | h_t)} \right] \right] \\
 &= \mathbb{E}_{d(M_k, \tau_{:H^+})} \left[\mathbb{E}_{q_\phi(y_t, z_t | h_t)} \left[\log \frac{p_\theta(\tau_{:H^+} | y_t, z_t) p_\theta(z_t | y_t) p(y_t)}{q_\phi(y_t | h_t) q_\phi(z_t | h_t, y_t)} \right] \right] \\
 &= \mathbb{E}_{d(M_k, \tau_{:H^+})} \left[\mathbb{E}_{q_\phi(y_t, z_t | h_t)} \left[\log p_\theta(\tau_{:H^+} | z_t) + \frac{p_\theta(z_t | y_t)}{q_\phi(z_t | h_t, y_t)} + \log \frac{p(y_t)}{q_\phi(y_t | h_t)} \right] \right]. \tag{8}
 \end{aligned}$$

527 Eq. (8) is equivalent to the ELBO objective in Eq. (3) without weighting coefficients. We assume
528 that the reconstruction $\tau_{:H^+}$ is conditionally independent of the subtask composition y_t given z_t .

529 **C Implementation Details**

530 **C.1 Reference Implementations**

531 **MAML, RL², and PEARL** To replicate the results of MAML [12], RL² [8], and PEARL [44] as
 532 reported in the Meta-World paper [69], we utilize the exact version² of the Garage repository [14]
 533 without any modifications. For their hyperparameters, please refer to Appendix D.7, D.8, and D.9 of
 534 the Meta-World paper. MAML is the only baseline that has the advantage to take gradient updates
 535 during the test.

536 **SDVT, SD, LDM, VariBAD** Task-inference-based methods (SDVT, SD, LDM [33], and VariBAD
 537 [72]) are adapted to the Meta-World benchmark based on the VariBAD’s implementation.³ We
 538 begin with VariBAD’s hyperparameter configuration for MuJoCo [56] Ant-goal and modify certain
 539 hyperparameters to accommodate the Meta-World benchmark (e.g., batch size, network capacity,
 540 etc.). In line with VariBAD and LDM, we train SDVT’s policy using PPO [47]. The GMVAE is
 541 based on an implementation⁴ that employs the Gumbel-Softmax reparameterization trick [22] when
 542 sampling the categorical subtask variable. The virtual training processes of SDVT and LDM require
 543 agent-environment interaction since they use states from the real environment. Therefore, the virtual
 544 training steps are also added when counting the total number of training steps. For a comprehensive
 545 list of SDVT hyperparameters, shared among SD, LDM, and VariBAD to ensure a fair comparison,
 546 please refer to Appendix C.3 and the supplementary material’s source code.

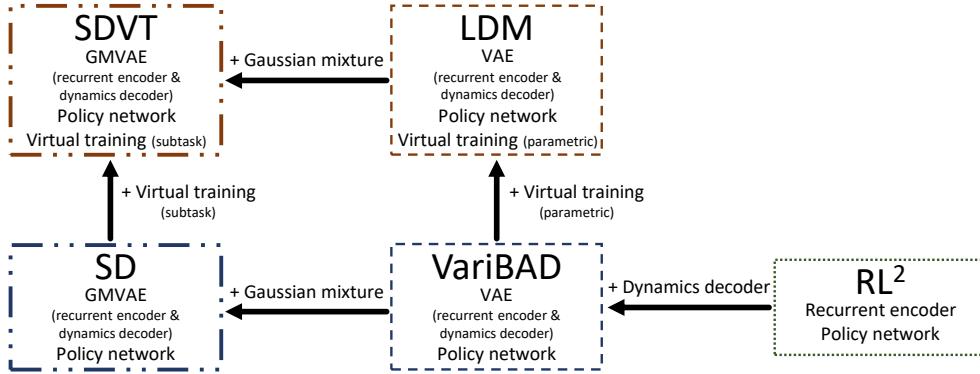


Figure 5: **Schematic overview of algorithms.** SDVT without a Gaussian mixture reduces to LDM, LDM without virtual training reduces to VariBAD, and VariBAD without a VAE decoder reduces to RL². MAML and PEARL use fully connected networks.

547 **C.2 Computational Complexity**

Table 2: **Computational complexity.** The total wall-clock time required to generate the results for the ML-10, averaged across all eight random seeds.

	SDVT	SDVT-LW	SD	SD-LW	RL ²	MAML	PEARL	VariBAD	LDM
Wall-clock time (hours)	142	140	138	135	192	17	258	126	131

548 Our experiments were conducted using an Nvidia TITAN Xp. Due to the considerably larger variety
 549 of tasks compared to standard meta-RL benchmarks, non-parametric benchmarks require significantly
 550 more computational resources and time. In Table 2, we detail the total wall-clock time expended by
 551 our methods and baselines during the training and evaluation of the ML-10 benchmark. This time
 552 includes the environmental interaction time for 250M training steps, roughly 100M steps dedicated to
 553 evaluations, and the time taken to train the neural networks. Despite incorporating the GMVAE and
 554 virtual training, our method’s computational demand does not substantially surpass that of VariBAD.

²<https://github.com/rlworkgroup/garage/pull/2287>

³<https://github.com/lmzintgraf/varibad>

⁴<https://github.com/jariasf/GMVAE>

Table 3: **Hyperparameters of SDVT and SD.** Hyperparameters of SDVT used for Meta-World ML-10 and ML-45 along with the notations in the manuscript and the argument names in the source code. SD shares the same hyperparameters except without those of virtual training. Different hyperparameters of our lightweight variant are denoted as (-LW).

Category	Description	Notation	Value (ML-10, 45)	Argument Name
General	Rollout episode horizon	H	500	max_episode_steps
	Number of rollout episodes	n_{roll}	10	max_rollouts_per_task
	Discount factor	γ	0.99	policy_gamma
	Number of parallel processes		10	num_processes
	Total Environment Steps		2.5e8, 4.0e8	num_frames
GMVAE	Optimizer		Adam	optimiser_vae
	Learning rate		$1e - 3$	lr_vae
	Subsample ELBO		100	vae_subsample_elbos
	Subsample decodes		100	vae_subsample_decodes
	Buffer size (meta-episodes)	$ \mathcal{B}_{\text{VAE}} $	1000	size_vae_buffer
	ELBO categorical coefficient	α_c	1.0	cat_loss_coeff
	ELBO categorical coefficient (-LW)	α_c	0.5	cat_loss_coeff
	ELBO Gaussian coefficient	α_g	1.0	gauss_loss_coeff
	ELBO reward coefficient	α_R	10	rew_loss_coeff
	ELBO transition coefficient	α_T	1000	state_loss_coeff
	ELBO occupancy coefficient	α_o	1.0	occ_loss_coeff
	ELBO occupancy coefficient (-LW)	α_o	0	occ_loss_coeff
	Gumbel softmax temperature		1	gumbel_temperature
	Number of updates per epoch		10, 20	num_vae_updates
	Subtask dimension	$K = N_{\text{train}}$	10, 45	vae_mixture_num
	Subtask dimension (-LW)	K	5	vae_mixture_num
	Gaussian dimension	$\dim(z)$	5, 10	latent_dim
Policy	Dropout rate	p_{drop}	0.7	dropout_rate
	Reward reconstruction objective		MSE	rew_pred_type
	State reconstruction objective		MSE	state_pred_type
	Algorithm	PPO		policy
	Optimiser	Adam		policy_optimiser
Virtual training	Learning rate	$7e - 4$		lr_policy
	Optimizer epsilon	$10e - 8$		policy_eps
	PPO update epochs	5		ppo_num_epochs
	Steps per policy update	$\frac{ \mathcal{B}_{\text{pol}} }{\text{num_processes}}$	5000	policy_num_steps
	Number of minibatches		10	ppo_num_minibatch
	PPO clipping parameter		0.1	ppo_clip_param
	GAE λ		0.9	policy_tau
	Initial standard deviation		1.0	policy_init_std
	Minimum standard deviation		0.5	policy_min_std
	Maximum standard deviation		1.5	policy_max_std
	Entropy coefficient		$1e - 3$	policy_entropy_coeff
	Initial virtual ratio		0.0	virtual_ratio
	Virtual ratio increment per step	Δp_v	$5e - 8, 2.5e - 8$	virtual_ratio_increment
	ELBO dispersion coefficient	α_d	10	ext_loss_coeff

Table 4: **Hyperparameters of VariBAD and LDM.** VariBAD and LDM employ identical hyperparameters for the general and policy categories of SDVT and SD, as shown in Table 3. The sole distinction lies in the structural variation resulting from the utilization of GMVAE and VAE. We choose the Gaussian dimension that yielded the most favorable outcomes among 5, 10, 15, and 20.

Category	Description	Notation	Value (ML-10, 45)	Argument Name
VAE	Optimizer		Adam	optimiser_vae
	Learning rate		$1e - 3$	lr_vae
	Subsample ELBO		100	vae_subsample_elbos
	Subsample decodes		100	vae_subsample_decodes
	Buffer size (meta-episodes)	$ \mathcal{B}_{\text{VAE}} $	1000	size_vae_buffer
	ELBO KL coefficient		0.1	kl_weight
	ELBO reward coefficient	α_R	10	rew_loss_coeff
	ELBO transition coefficient	α_T	1000	state_loss_coeff
	Number of updates per epoch		10, 20	num_vae_updates
	Gaussian dimension	$\dim(z)$	5, 10	latent_dim
	Dropout rate (VariBAD)	p_{drop}	0.0	dropout_rate
	Dropout rate (LDM)	p_{drop}	0.7	dropout_rate
	Reward reconstruction objective		MSE	rew_pred_type
	State reconstruction objective		MSE	state_pred_type

556 C.4 Network Architecture

557 The network architectures comprising our method are described in Table 5. Prior to being input into
558 the encoder or decoder, all state, action, and reward inputs pass through embedding networks. The
559 values of $K = \dim(y)$ and $\dim(z)$ differ for the ML-10 and ML-45, as indicated in Table 3.

Table 5: **Network architecture.** Details of the network architecture composing the GMVAE. The numbers in the Layers column represent the dimension of the hidden layers and the output.

Network	Notation	Architecture	Layers	Activations (last layer)
State embedding		MLP	[32]	(Tanh)
Reward embedding		MLP	[16]	(Tanh)
Action embedding		MLP	[16]	(Tanh)
Recurrent encoder	q_{ϕ_h}	GRU	[256]	(None)
Categorical parameter	ω_{ϕ_y}	MLP	[512, 512, K]	ReLU (Softmax)
Gaussian parameters	$\mu_{\phi_z}, \sigma_{\phi_z}^2$	MLP	[512, 512, $\dim(z)$]	ReLU (None, SoftPlus)
Gaussian regularization parameters	$\mu_{\theta_z}, \sigma_{\theta_z}^2$	MLP	[$\dim(z)$]	(None, SoftPlus)
Latent context dispersion	p_{θ_h}	MLP	[256, 256, 256]	ReLU (None)
Reward decoder	p_{θ_R}	MLP	[64, 64, 32, 1]	ReLU (None)
Transition decoder	p_{θ_T}	MLP	[64, 64, 32, 40]	ReLU (None)
Policy	π_{ψ}	MLP	[256, 256, 4]	Tanh (Tanh)

560 C.5 Aggregation method for the Success Rate

561 The main results in the Meta-World paper (Table 1 and Figure 6 of Yu et al. [69]) present the average
562 of the *maximum* success rate for each task, diverging from the raw scores in their Figure 17 and
563 18. For instance, if an agent achieves a 90% success rate on “Door-close” in one evaluation during
564 training and scores 10% in all other evaluations at different times, the reported success rate for
565 “Door-close” is 90%. The mean of these maximum scores across tasks is reported as the aggregated
566 success rate.

567 This unconventional aggregation method does not accurately represent the meta-RL objective, which
568 aims to train a single agent capable of solving multiple tasks. Meta-World calculates success rates for
569 various tasks at distinct time points during training, even though the agent might specialize in different
570 tasks at various stages. As a result, the agent may excel at specific tasks by chance. Consequently,
571 evaluating the agent more frequently could yield higher *maximum* scores. We report the conventional
572 final performance, which better reflects the meta-RL objective.

573 **D Detailed Experimental Results**

574 **D.1 Performance on Individual Tasks**

Table 6: Meta-World V2 ML-10 success rate. We report the final success rates (%) of baselines and our methods for training tasks and test tasks of the Meta-World ML-10 benchmark. All results are reported as the mean success rate \pm 95% confidence interval of 8 seeds at 250M steps.

Index. Task	SDVT	SDVT-LW	SD	SD-LW	RL ²	MAML	PEARL	VariBAD	LDM
1. Reach	50.0 \pm 11.0	22.0 \pm 3.5	53.8 \pm 11.5	56.2 \pm 9.1	27.5 \pm 19.8	50.0 \pm 34.6	70.0\pm11.9	28.5 \pm 5.0	28.5 \pm 4.1
2. Push	57.5 \pm 18.6	25.2 \pm 16.5	61.2 \pm 7.3	68.8\pm10.6	60.0 \pm 13.0	0.0 \pm 0.0	0.0 \pm 0.0	25.8 \pm 12.9	25.5 \pm 12.8
3. Pick-place	47.5 \pm 9.6	39.0 \pm 5.4	52.5 \pm 7.6	55.0\pm13.9	42.5 \pm 21.0	0.0 \pm 0.0	0.0 \pm 0.0	24.8 \pm 13.3	27.2 \pm 13.3
4. Door-open	100.0\pm0.0	60.5 \pm 31.9	100.0\pm0.0	63.8 \pm 32.5	93.8 \pm 9.1	79.5 \pm 22.1	10.2 \pm 11.9	86.8 \pm 22.7	62.5 \pm 33.5
5. Drawer-close	100.0\pm0.0	100.0\pm0.0	100.0\pm0.0	100.0\pm0.0	98.8 \pm 2.3	99.5 \pm 1.1	86.8 \pm 5.6	100.0\pm0.0	100.0\pm0.0
6. Button-press	98.8 \pm 2.3	98.8 \pm 1.8	96.2 \pm 6.9	100.0\pm0.0	87.5 \pm 8.3	93.0 \pm 8.3	46.0 \pm 11.2	98.5 \pm 0.9	99.5 \pm 0.9
7. Peg-insert-side	31.2 \pm 13.2	24.5 \pm 10.2	50.0\pm19.9	36.2 \pm 11.5	45.0 \pm 21.6	0.0 \pm 0.0	0.0 \pm 0.0	19.8 \pm 14.0	21.8 \pm 15.3
8. Window-open	100.0\pm0.0	100.0\pm0.0	98.8 \pm 2.3	100.0\pm0.0	100.0\pm0.0	100.0\pm0.0	19.2 \pm 11.5	100.0\pm0.0	100.0\pm0.0
9. Sweep	96.2\pm4.8	72.2 \pm 29.0	83.8 \pm 22.2	83.8 \pm 15.1	75.0 \pm 17.3	0.0 \pm 0.0	0.0 \pm 0.0	60.2 \pm 32.4	61.2 \pm 32.9
10. Basketball	91.2\pm6.4	78.8 \pm 9.8	73.8 \pm 14.3	91.2 \pm 8.8	43.8 \pm 16.6	0.0 \pm 0.0	0.0 \pm 0.0	37.2 \pm 26.0	40.5 \pm 28.9
Train mean	77.2\pm3.0	62.1 \pm 4.1	77.0 \pm 5.9	75.5 \pm 5.5	67.4 \pm 4.4	42.2 \pm 4.5	23.2 \pm 1.9	58.2 \pm 8.9	56.7 \pm 12.3
11. Drawer-open	65.0\pm19.9	30.5 \pm 12.9	48.8 \pm 23.4	45.0 \pm 24.0	2.2 \pm 1.9	15.8 \pm 19.3	1.5 \pm 1.1	12.8 \pm 12.8	21.8 \pm 12.0
12. Door-close	7.5 \pm 9.0	81.2\pm19.0	33.8 \pm 25.4	18.8 \pm 24.1	8.2 \pm 7.3	3.2 \pm 6.0	1.2 \pm 1.5	27.0 \pm 21.8	30.2 \pm 28.2
13. Shelf-place	0.0 \pm 0.0	1.0\pm1.2	0.0 \pm 0.0	0.0 \pm 0.0	0.2 \pm 0.2	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
14. Sweep-into	90.0\pm8.5	51.2 \pm 18.9	71.2 \pm 14.9	55.0 \pm 21.6	64.5 \pm 8.3	0.0 \pm 0.0	0.8 \pm 1.0	30.5 \pm 22.0	46.5 \pm 21.8
15. Lever-pull	1.2 \pm 2.3	3.2 \pm 2.7	0.0 \pm 0.0	12.5\pm10.2	0.5 \pm 0.8	0.5 \pm 0.9	0.5 \pm 0.9	0.2 \pm 0.5	0.5 \pm 0.9
Test mean	32.8 \pm 3.9	33.4\pm5.0	30.8 \pm 7.7	26.2 \pm 8.7	15.1 \pm 2.7	3.9 \pm 3.7	0.8 \pm 0.5	14.1 \pm 6.1	19.8 \pm 6.0

Table 7: Meta-World V2 ML-10 return. We report the final return of baselines and our methods at the last rollout episode on the Meta-World ML-10 analogous to the success rate in Table 6. Results are reported as the mean \pm 95% confidence intervals of 8 seeds at 250M steps.

Index. Task	SDVT	SDVT-LW	SD	SD-LW	RL ²	MAML	PEARL	VariBAD	LDM
1. Reach	3565 \pm 132	3681 \pm 115	3708 \pm 228	3806 \pm 220	2291 \pm 276	3899\pm560	3841 \pm 554	3829 \pm 153	3740 \pm 59
2. Push	3923\pm339	3493 \pm 469	3636 \pm 628	3579 \pm 870	815 \pm 133	17 \pm 1	15 \pm 2	2609 \pm 990	2470 \pm 1220
3. Pick-place	2257\pm239	2241 \pm 300	2131 \pm 441	2156 \pm 593	501 \pm 137	6 \pm 0	5 \pm 1	1395 \pm 759	1561 \pm 861
4. Door-open	4441 \pm 37	3171 \pm 825	4447\pm106	3530 \pm 831	1249 \pm 115	3561 \pm 753	1064 \pm 254	4051 \pm 780	3249 \pm 1134
5. Drawer-close	4813 \pm 71	4850 \pm 10	4850 \pm 13	4854\pm9	2193 \pm 336	4804 \pm 55	4081 \pm 314	4841 \pm 18	4837 \pm 18
6. Button-press	3392 \pm 59	3437\pm133	3257 \pm 265	3423 \pm 209	1189 \pm 151	2574 \pm 274	956 \pm 300	3246 \pm 161	3252 \pm 140
7. Peg-insert-side	1903 \pm 430	1919 \pm 348	2634\pm588	2225 \pm 592	682 \pm 157	7 \pm 0	7 \pm 1	1148 \pm 761	1237 \pm 825
8. Window-open	4486\pm39	4371 \pm 77	4349 \pm 165	4398 \pm 79	1230 \pm 112	3271 \pm 732	784 \pm 390	4483 \pm 32	4458 \pm 42
9. Sweep	4235\pm202	3617 \pm 682	3823 \pm 702	3557 \pm 634	906 \pm 223	79 \pm 38	51 \pm 14	2807 \pm 1260	2789 \pm 1333
10. Basketball	3549 \pm 292	3759\pm307	3466 \pm 511	3726 \pm 331	532 \pm 116	7 \pm 1	6 \pm 2	2141 \pm 1119	2035 \pm 1256
Train mean	3656\pm62	3454 \pm 137	3630 \pm 241	3525 \pm 297	1159 \pm 83	1822 \pm 136	1081 \pm 77	3055 \pm 466	2963 \pm 626
11. Drawer-open	2558\pm165	2280 \pm 271	2305 \pm 298	2310 \pm 280	1649 \pm 100	1737 \pm 270	1130 \pm 87	2336 \pm 340	2389 \pm 386
12. Door-close	566 \pm 335	3030\pm857	1068 \pm 740	638 \pm 670	568 \pm 137	263 \pm 278	321 \pm 221	1121 \pm 689	1343 \pm 1214
13. Shelf-place	501 \pm 77	550 \pm 138	494 \pm 102	434 \pm 154	578\pm70	0 \pm 0	0 \pm 0	211 \pm 144	328 \pm 204
14. Sweep-into	2221\pm764	1455 \pm 615	1359 \pm 632	1532 \pm 721	490 \pm 46	41 \pm 3	48 \pm 33	658 \pm 358	1509 \pm 889
15. Lever-pull	279 \pm 38	318 \pm 45	335\pm45	300 \pm 48	291 \pm 66	156 \pm 49	200 \pm 58	271 \pm 26	260 \pm 48
Test mean	1225 \pm 160	1527\pm214	1112 \pm 190	1043 \pm 234	715 \pm 33	439 \pm 78	340 \pm 54	919 \pm 143	1166 \pm 264

Table 8: **Meta-World V2 ML-45 success rate.** We report the final success rates (%) of baselines and our methods for training tasks and test tasks of the Meta-World ML-45 benchmark. All results are reported as the mean success rate \pm 95% confidence interval of 8 seeds at 400M steps.

Index. Task	SDVT	SDVT-LW	SD	SD-LW	RL ²	MAML	PEARL	VariBAD	LDM
1. Assembly	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	2.2\pm0.4	0.0 \pm 0.0	0.0 \pm 0.1	0.0 \pm 0.0	0.0 \pm 0.0
2. Basketball	24.8\pm17.7	3.0 \pm 3.6	23.8 \pm 19.2	24.2 \pm 13.6	21.0 \pm 9.8	0.0 \pm 0.0	0.0 \pm 0.0	19.8 \pm 13.7	10.2 \pm 8.9
3. Button-press -topdown	95.0 \pm 7.6	94.5 \pm 0.6	100.0\pm0.0	98.0 \pm 1.7	98.8 \pm 1.0	78.2 \pm 21.2	7.5 \pm 8.6	99.5 \pm 0.6	82.5 \pm 18.7
4. Button-press -topdown-wall	94.0 \pm 7.9	96.5 \pm 0.6	99.0 \pm 1.4	99.5 \pm 0.6	98.5 \pm 1.5	78.2 \pm 21.6	9.5 \pm 10.6	99.8\pm0.5	86.2 \pm 15.6
5. Button-press	93.2 \pm 3.3	96.5 \pm 0.6	89.0 \pm 8.1	99.0\pm0.7	84.2 \pm 3.2	95.5 \pm 1.9	26.8 \pm 11.3	98.0 \pm 1.0	93.5 \pm 8.0
6. Button-press -wall	86.0\pm11.6	68.0 \pm 4.0	80.5 \pm 8.7	75.5 \pm 9.4	67.0 \pm 4.8	71.2 \pm 10.6	20.2 \pm 8.3	74.5 \pm 8.3	56.0 \pm 11.3
7. Coffee-button	100.0\pm0.0	100.0\pm0.0	100.0\pm0.0	97.5 \pm 2.3	97.5 \pm 1.6	92.0 \pm 4.9	2.8 \pm 1.3	78.0 \pm 14.7	100.0\pm0.0
8. Coffee-pull	37.2 \pm 19.1	33.2 \pm 12.2	45.5 \pm 23.4	46.5 \pm 22.0	65.0\pm4.3	0.5 \pm 0.3	0.0 \pm 0.1	52.5 \pm 10.5	48.8 \pm 12.9
9. Coffee-push	64.8 \pm 6.5	37.0 \pm 15.8	56.8 \pm 2.7	67.8 \pm 12.6	60.8 \pm 9.0	22.0 \pm 10.6	3.0 \pm 2.4	76.5 \pm 10.8	81.8\pm5.8
10. Dial-turn	87.2\pm4.4	86.2 \pm 2.9	82.8 \pm 7.9	67.2 \pm 11.9	14.5 \pm 8.3	81.2 \pm 8.8	9.8 \pm 7.7	44.0 \pm 13.9	60.8 \pm 5.9
11. Disassemble	59.5 \pm 32.0	32.0 \pm 14.7	87.8\pm6.2	74.0 \pm 10.9	23.2 \pm 27.1	0.0 \pm 0.1	0.2 \pm 0.2	44.5 \pm 21.8	65.5 \pm 18.5
12. Door-close	100.0\pm0.0	100.0\pm0.0	99.5 \pm 0.6	97.0 \pm 3.6	98.2 \pm 1.2	100.0\pm0.0	52.2 \pm 21.2	100.0\pm0.0	99.2 \pm 1.0
13. Door-open	56.2 \pm 31.2	56.0 \pm 24.3	99.8\pm0.5	64.8 \pm 26.9	54.0 \pm 23.8	6.8 \pm 11.9	0.5 \pm 0.3	98.0 \pm 1.4	96.2 \pm 3.6
14. Drawer-close	100.0\pm0.0	99.0 \pm 1.2	97.2 \pm 2.2	100.0\pm0.0	98.5 \pm 0.5	100.0 \pm 0.1	96.5 \pm 2.9	99.2 \pm 0.7	96.5 \pm 4.0
15. Drawer-open	96.5 \pm 1.1	83.8 \pm 16.3	96.0 \pm 4.8	87.2 \pm 13.8	83.5 \pm 8.4	0.0 \pm 0.0	2.5 \pm 1.3	100.0\pm0.0	93.2 \pm 6.6
16. Faucet-open	99.5 \pm 0.6	82.5 \pm 21.0	99.2 \pm 1.0	100.0\pm0.0	86.8 \pm 9.7	70.5 \pm 19.1	28.8 \pm 11.7	95.5 \pm 1.7	98.2 \pm 1.5
17. Faucet-close	97.0 \pm 5.5	77.0 \pm 18.3	99.2 \pm 0.7	97.0 \pm 2.3	56.8 \pm 16.8	52.8 \pm 22.5	15.0 \pm 6.3	99.5\pm0.6	99.2 \pm 1.0
18. Hammer	2.0 \pm 3.7	0.0 \pm 0.0	0.0 \pm 0.0	1.5 \pm 1.8	12.8\pm15.4	6.8 \pm 5.8	2.2 \pm 1.3	0.0 \pm 0.0	0.0 \pm 0.0
19. Handle-press -side	97.8 \pm 4.1	100.0\pm0.0	100.0\pm0.0	100.0\pm0.0	99.8 \pm 0.1	83.2 \pm 12.2	12.8 \pm 11.8	100.0\pm0.0	99.8 \pm 0.5
20. Handle-press	98.5 \pm 1.1	99.0 \pm 1.2	100.0\pm0.0	100.0\pm0.0	99.8 \pm 0.1	72.0 \pm 20.2	66.8 \pm 19.8	100.0\pm0.0	100.0\pm0.0
21. Handle-pull -side	40.0 \pm 22.3	5.2 \pm 5.2	50.5 \pm 25.2	39.5 \pm 24.1	83.2\pm4.3	13.8 \pm 15.5	0.8 \pm 0.7	41.2 \pm 23.3	31.8 \pm 28.5
22. Handle-pull	0.0 \pm 0.0	1.0 \pm 1.2	35.2 \pm 31.5	3.8 \pm 2.6	0.2 \pm 0.2	49.0\pm18.9	0.8 \pm 0.4	1.0 \pm 0.7	1.0 \pm 1.0
23. Lever-pull	10.0 \pm 9.8	16.0 \pm 4.4	11.5 \pm 6.8	23.2\pm16.6	1.0 \pm 0.8	3.2 \pm 3.9	0.0 \pm 0.0	20.0 \pm 11.3	6.0 \pm 3.9
24. Peg-insert -side	1.0 \pm 0.7	0.0 \pm 0.0	0.8 \pm 0.7	0.5 \pm 0.6	19.5\pm5.8	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	2.2 \pm 2.7
25. Pick-place -wall	30.0 \pm 16.7	29.5 \pm 9.8	48.5\pm3.7	41.2 \pm 11.7	42.5 \pm 4.0	6.5 \pm 6.0	0.0 \pm 0.0	35.5 \pm 4.6	31.2 \pm 15.9
26. Pick-out -of-hole	26.8 \pm 16.4	23.2 \pm 12.0	53.0\pm7.0	43.8 \pm 8.4	5.5 \pm 2.7	0.0 \pm 0.0	0.0 \pm 0.0	32.0 \pm 20.8	34.0 \pm 11.4
27. Reach	31.5 \pm 6.7	13.5 \pm 4.0	36.0 \pm 8.7	25.5 \pm 6.7	45.8\pm8.4	18.8 \pm 10.6	13.5 \pm 7.2	27.0 \pm 4.7	28.5 \pm 3.2
28. Push-back	40.8 \pm 21.3	31.0 \pm 11.9	76.0 \pm 5.3	44.5 \pm 9.3	77.8\pm4.6	0.0 \pm 0.0	0.0 \pm 0.0	41.5 \pm 24.1	49.5 \pm 10.5
29. Push	50.5 \pm 5.7	62.8\pm5.1	36.2 \pm 14.3	29.2 \pm 18.3	58.5 \pm 1.6	16.8 \pm 14.9	1.8 \pm 1.4	45.0 \pm 9.5	35.2 \pm 11.7
30. Pick-place	28.2 \pm 13.5	26.2 \pm 6.5	57.5\pm9.1	36.8 \pm 3.9	43.5 \pm 1.8	4.8 \pm 5.3	0.0 \pm 0.0	34.8 \pm 3.9	44.0 \pm 5.2
31. Plate-slide	75.8\pm5.7	56.5 \pm 6.9	58.2 \pm 10.8	60.2 \pm 10.4	20.5 \pm 14.4	17.0 \pm 15.2	9.5 \pm 10.6	55.2 \pm 6.8	57.0 \pm 7.3
32. Plate-slide -side	81.8 \pm 0.8	1.0 \pm 0.7	97.0\pm3.6	87.0 \pm 3.2	70.8 \pm 5.1	2.0 \pm 3.6	0.2 \pm 0.3	89.5 \pm 13.6	59.5 \pm 24.2
33. Plate-slide -back	82.0 \pm 6.2	86.0 \pm 9.7	83.5 \pm 4.2	76.5 \pm 10.8	90.0\pm3.9	0.0 \pm 0.0	0.8 \pm 0.8	84.0 \pm 6.1	80.2 \pm 7.9
34. Plate-slide -back-side	71.8 \pm 5.6	59.0 \pm 4.5	60.8 \pm 11.2	73.5 \pm 10.8	81.5\pm1.9	5.5 \pm 5.7	0.5 \pm 0.6	73.5 \pm 2.4	80.8 \pm 8.0
35. Peg-unplug -side	40.0 \pm 17.2	50.8 \pm 12.0	54.0 \pm 5.1	47.8 \pm 17.8	54.0 \pm 3.7	13.0 \pm 3.4	3.5 \pm 2.1	60.5\pm5.8	48.8 \pm 9.1
36. Soccer	29.8 \pm 7.2	25.5 \pm 8.8	16.0 \pm 2.7	24.5 \pm 16.9	37.8\pm4.7	14.5 \pm 12.7	2.5 \pm 1.8	20.2 \pm 5.4	15.2 \pm 2.5
37. Stick-push	0.0 \pm 0.0	47.5 \pm 32.9	0.0 \pm 0.0	20.5 \pm 24.6	88.2\pm3.4	0.0 \pm 0.1	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
38. Stick-pull	0.0 \pm 0.0	15.0\pm10.4	0.0 \pm 0.0	12.0 \pm 14.4	4.5 \pm 3.6	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0	0.0 \pm 0.0
39. Push-wall	51.0 \pm 4.9	50.0 \pm 25.4	62.8 \pm 21.7	27.5 \pm 20.8	72.5\pm5.4	18.2 \pm 16.5	0.2 \pm 0.2	63.8 \pm 15.7	44.0 \pm 20.5
40. Reach-wall	16.2 \pm 5.9	7.0 \pm 4.9	27.0 \pm 8.1	28.0 \pm 7.3	65.0\pm5.6	35.2 \pm 20.2	15.0 \pm 6.9	18.0 \pm 3.6	13.8 \pm 9.1
41. Shelf-place	0.0 \pm 0.0	0.5 \pm 0.6	0.8 \pm 0.7	0.0 \pm 0.0	3.8 \pm 2.9	0.0 \pm 0.0	0.0 \pm 0.0	4.5\pm4.6	0.8 \pm 1.0
42. Sweep-into	92.8 \pm 3.2	91.0 \pm 2.9	90.5 \pm 5.8	90.2 \pm 8.6	87.5 \pm 1.5	23.8 \pm 21.3	4.8 \pm 3.2	97.8 \pm 2.6	98.5\pm1.3
43. Sweep	11.8 \pm 10.1	27.0 \pm 17.4	33.0 \pm 17.7	20.0 \pm 23.2	41.5\pm21.1	0.0 \pm 0.0	0.0 \pm 0.0	40.5 \pm 28.1	5.5 \pm 6.5
44. Window-open	99.5 \pm 0.6	100.0\pm0.0	100.0\pm0.0	100.0\pm0.0	97.8 \pm 1.0	91.5 \pm 6.6	27.5 \pm 11.4	98.2 \pm 1.6	98.8 \pm 2.3
45. Window-close	100.0\pm0.0	99.0 \pm 1.2	100.0\pm0.0	100.0\pm0.0	95.2 \pm 3.3	97.0 \pm 1.9	27.2 \pm 11.0	100.0\pm0.0	99.8 \pm 0.5
Train mean	55.6 \pm 4.2	50.4 \pm 4.1	61.0\pm1.7	56.7 \pm 1.5	58.0 \pm 0.4	32.0 \pm 1.4	10.3 \pm 2.4	57.0 \pm 1.2	54.1 \pm 0.9
46. Bin-picking	0.0 \pm 0.0	1.5 \pm 1.8	2.2\pm1.1	1.2 \pm 1.0	1.2 \pm 0.9	0.0 \pm 0.0	0.0 \pm 0.0	1.8 \pm 1.1	0.5 \pm 0.9
47. Box-close	0.5 \pm 0.6	0.0 \pm 0.0	0.5 \pm 0.6	0.0 \pm 0.0	0.5 \pm 0.3	0.0 \pm 0.0	0.0 \pm 0.0	5.5\pm4.6	0.0 \pm 0.0
48. Hand-insert	2.2 \pm 2.3	3.8 \pm 3.1	1.0 \pm 1.2	1.5 \pm 1.8	5.2 \pm 3.0	23.0\pm15.8	0.0 \pm 0.0	1.5 \pm 1.1	0.8 \pm 1.0
49. Door-lock	59.2 \pm 9.1	74.8\pm7.5	49.8 \pm 15.5	51.5 \pm 9.5	14.0 \pm 8.2	11.2 \pm 10.5	8.2 \pm 10.0	37.5 \pm 20.3	59.2 \pm 9.6
50. Door-unlock	78.5\pm7.9	75.8 \pm 9.1	61.8 \pm 10.2	72.5 \pm 9.5	38.0 \pm 11.3	64.5 \pm 19.2	25.0 \pm 19.2	64.2 \pm 4.3	63.2 \pm 6.9
Test mean	28.1 \pm 3.2	31.2\pm1.2	23.0 \pm 5.1	25.4 \pm 2.9	11.8 \pm 3.2	19.8 \pm 6.3	6.7 \pm 3.3	22.1 \pm 3.5	24.8 \pm 2.9

Table 9: **Meta-World V2 ML-45 return.** We report the final return of baselines and our methods at the last rollout episode on the Meta-World ML-45 analogous to the success rate in Table 8. Results are reported as the mean \pm 95% confidence intervals of 8 seeds at 400M steps.

Index. Task	SDVT	SDVT-LW	SD	SD-LW	RL ²	MAML	PEARL	VariBAD	LDM
1. Assembly	1813 \pm 485	2573\pm117	2414 \pm 233	2285 \pm 195	1202 \pm 19	268 \pm 70	208 \pm 9	2434 \pm 161	2572 \pm 56
2. Basketball	1929 \pm 688	1443 \pm 524	2215 \pm 148	2581\pm129	489 \pm 38	100 \pm 90	6 \pm 2	1311 \pm 651	2227 \pm 236
3. Button-press-topdown	3120 \pm 447	3297 \pm 328	3580\pm87	3343 \pm 176	1344 \pm 14	2660 \pm 633	543 \pm 273	3348 \pm 227	3208 \pm 354
4. Button-press-topdown-wall	3217 \pm 384	3386 \pm 246	3614\pm110	3392 \pm 185	1359 \pm 18	2664 \pm 638	604 \pm 314	3337 \pm 279	3218 \pm 319
5. Button-press	3100\pm270	2977 \pm 199	2833 \pm 208	2938 \pm 193	1165 \pm 29	2942 \pm 178	434 \pm 195	2878 \pm 74	2745 \pm 205
6. Button-press-wall	3401\pm220	2916 \pm 94	3148 \pm 157	3165 \pm 92	1378 \pm 17	2589 \pm 341	327 \pm 183	3040 \pm 136	2991 \pm 108
7. Coffee-button	1605 \pm 678	2196 \pm 430	2517 \pm 296	2420 \pm 303	1348 \pm 40	3342\pm157	80 \pm 20	2253 \pm 303	2610 \pm 271
8. Coffee-pull	1032 \pm 307	916 \pm 323	1010 \pm 338	1305\pm395	572 \pm 35	38 \pm 7	21 \pm 5	1292 \pm 260	1244 \pm 220
9. Coffee-push	1424 \pm 443	889 \pm 455	1254 \pm 137	1687 \pm 497	504 \pm 41	111 \pm 47	20 \pm 10	1901 \pm 413	2175\pm491
10. Dial-turn	1821 \pm 478	1704 \pm 520	1835\pm414	1825 \pm 484	944 \pm 70	1553 \pm 254	449 \pm 407	680 \pm 130	1101 \pm 571
11. Disassemble	2451 \pm 1226	897 \pm 350	3387\pm382	2840 \pm 485	2642 \pm 1397	130 \pm 58	150 \pm 25	1564 \pm 749	2268 \pm 688
12. Door-close	4178 \pm 204	4404\pm22	4375 \pm 17	4251 \pm 93	1231 \pm 20	4265 \pm 158	2430 \pm 965	4181 \pm 105	4067 \pm 174
13. Door-open	3086 \pm 761	2215 \pm 430	4134\pm151	2833 \pm 517	1312 \pm 216	1095 \pm 316	469 \pm 142	3741 \pm 333	3857 \pm 325
14. Drawer-close	4310 \pm 274	4544 \pm 54	4392 \pm 307	4680 \pm 80	1626 \pm 25	4725\pm28	4539 \pm 225	4515 \pm 112	4506 \pm 206
15. Drawer-open	3964 \pm 96	3624 \pm 334	4265\pm106	4035 \pm 221	2169 \pm 118	2259 \pm 235	1288 \pm 367	4180 \pm 114	3876 \pm 270
16. Faucet-open	4572 \pm 68	4057 \pm 692	4582 \pm 81	4674\pm21	1925 \pm 193	3575 \pm 540	1810 \pm 497	4338 \pm 155	4479 \pm 118
17. Faucet-close	4166 \pm 351	3656 \pm 471	4457\pm194	4300 \pm 95	1662 \pm 115	3298 \pm 561	1470 \pm 414	4267 \pm 57	4401 \pm 196
18. Hammer	466 \pm 60	465 \pm 17	469 \pm 5	459 \pm 16	738\pm233	649 \pm 135	420 \pm 20	476 \pm 4	384 \pm 115
19. Handle-press-side	4676 \pm 192	4722 \pm 14	4767 \pm 38	4797\pm17	2276 \pm 43	3696 \pm 522	652 \pm 524	4772 \pm 19	4737 \pm 51
20. Handle-press	4601 \pm 74	4355 \pm 183	4480 \pm 147	4789\pm14	2239 \pm 86	2902 \pm 791	3036 \pm 958	4728 \pm 22	4747 \pm 55
21. Handle-pull-side	853 \pm 523	83 \pm 30	1507\pm832	907 \pm 731	565 \pm 43	199 \pm 300	12 \pm 5	991 \pm 576	917 \pm 826
22. Handle-pull	1469 \pm 153	1491 \pm 110	2260\pm666	1541 \pm 100	2048 \pm 86	2001 \pm 597	29 \pm 19	1549 \pm 57	1612 \pm 170
23. Lever-pull	446 \pm 84	496 \pm 74	483 \pm 39	500\pm106	305 \pm 21	264 \pm 60	119 \pm 38	457 \pm 32	345 \pm 27
24. Peg-insert-side	767 \pm 408	757 \pm 48	1106\pm186	1067 \pm 192	640 \pm 77	33 \pm 24	7 \pm 2	778 \pm 101	962 \pm 232
25. Pick-place-wall	1647 \pm 893	2011 \pm 717	2765\pm64	2570 \pm 543	491 \pm 22	389 \pm 359	0 \pm 0	2042 \pm 501	1873 \pm 854
26. Pick-out-of-hole	809 \pm 435	1015 \pm 365	1319 \pm 183	1385\pm310	270 \pm 26	20 \pm 1	8 \pm 1	1129 \pm 261	1247 \pm 392
27. Reach	2682 \pm 225	2448 \pm 313	3048 \pm 238	3116\pm182	2273 \pm 49	2447 \pm 829	2466 \pm 1079	2843 \pm 147	3085 \pm 183
28. Push-back	674 \pm 362	514 \pm 154	1438\pm156	1218 \pm 334	433 \pm 42	11 \pm 7	8 \pm 3	871 \pm 372	1404 \pm 324
29. Push	3130 \pm 290	3139 \pm 655	2961 \pm 253	2581 \pm 530	945 \pm 32	634 \pm 555	34 \pm 13	3427 \pm 185	3524\pm200
30. Pick-place	1221 \pm 650	1582 \pm 433	2145\pm159	1992 \pm 122	576 \pm 23	247 \pm 222	5 \pm 1	2034 \pm 48	2128 \pm 144
31. Plate-slide	3526\pm161	2950 \pm 232	3022 \pm 196	2828 \pm 303	2863 \pm 1349	1043 \pm 789	561 \pm 490	2858 \pm 251	2801 \pm 231
32. Plate-slide-side	3058 \pm 289	1362 \pm 37	3471\pm258	2761 \pm 317	949 \pm 20	322 \pm 249	56 \pm 35	3107 \pm 167	2351 \pm 462
33. Plate-slide-back	3783 \pm 204	3790 \pm 239	3844 \pm 128	3754 \pm 272	1544 \pm 10	581 \pm 364	655 \pm 259	3900\pm155	3720 \pm 325
34. Plate-slide-back-side	3732 \pm 91	3764 \pm 149	3755 \pm 256	3987 \pm 228	1561 \pm 7	676 \pm 420	133 \pm 56	4099\pm72	4094 \pm 130
35. Peg-unplug-side	982 \pm 390	971 \pm 331	1187 \pm 111	1079 \pm 464	401 \pm 36	62 \pm 26	45 \pm 28	1523\pm142	1087 \pm 282
36. Soccer	1374 \pm 169	1566\pm207	1311 \pm 129	1348 \pm 332	671 \pm 31	439 \pm 359	98 \pm 60	1399 \pm 90	1486 \pm 82
37. Stick-push	47 \pm 70	786 \pm 533	11 \pm 2	291 \pm 333	955\pm14	128 \pm 212	6 \pm 1	17 \pm 2	10 \pm 4
38. Stick-pull	25 \pm 23	905\pm616	14 \pm 3	600 \pm 702	730 \pm 46	39 \pm 52	7 \pm 1	19 \pm 3	11 \pm 4
39. Push-wall	2547 \pm 426	2197 \pm 668	3291\pm421	2496 \pm 639	877 \pm 85	597 \pm 532	20 \pm 7	2813 \pm 716	2744 \pm 549
40. Reach-wall	1871 \pm 157	1889 \pm 412	2447 \pm 462	2741 \pm 456	2303 \pm 89	2802\pm1052	1741 \pm 641	2102 \pm 128	1836 \pm 569
41. Shelf-place	608 \pm 308	657 \pm 195	920 \pm 55	1056\pm165	623 \pm 61	9 \pm 8	0 \pm 0	795 \pm 186	725 \pm 167
42. Sweep-into	3755 \pm 333	3847 \pm 400	3869 \pm 214	3584 \pm 495	793 \pm 67	623 \pm 541	148 \pm 82	3750 \pm 159	4151\pm185
43. Sweep	1148 \pm 98	1433 \pm 477	1905\pm449	1415 \pm 623	872 \pm 230	160 \pm 82	36 \pm 9	1865 \pm 799	1219 \pm 325
44. Window-open	3789 \pm 249	4004 \pm 190	4138 \pm 56	4121 \pm 53	1163 \pm 37	2463 \pm 1004	705 \pm 184	4177\pm123	3986 \pm 154
45. Window-close	4201 \pm 166	4333 \pm 135	4293 \pm 209	4496\pm29	1152 \pm 109	3215 \pm 425	990 \pm 417	4346 \pm 111	4450 \pm 33
Train mean	2379 \pm 214	2294 \pm 202	2672\pm79	2578 \pm 64	1411 \pm 22	1388 \pm 104	597 \pm 121	2492 \pm 47	2515 \pm 67
46. Bin-picking	114 \pm 44	113 \pm 56	322\pm70	140 \pm 72	265 \pm 73	37 \pm 6	12 \pm 5	127 \pm 27	90 \pm 29
47. Box-close	190 \pm 49	128 \pm 2	143 \pm 5	137 \pm 9	820\pm159	215 \pm 26	176 \pm 40	243 \pm 83	151 \pm 10
48. Hand-insert	299 \pm 64	465 \pm 101	386 \pm 25	392 \pm 64	467\pm48	337 \pm 188	6 \pm 2	377 \pm 61	398 \pm 39
49. Door-lock	1694 \pm 302	1973\pm160	1359 \pm 262	1427 \pm 289	995 \pm 303	821 \pm 243	960 \pm 149	1358 \pm 445	1573 \pm 168
50. Door-unlock	1898\pm158	1792 \pm 225	1721 \pm 132	1870 \pm 245	769 \pm 108	1880 \pm 426	1374 \pm 537	1707 \pm 324	1627 \pm 169
Test mean	839 \pm 74	894\pm27	786 \pm 69	793 \pm 49	663 \pm 100	658 \pm 96	506 \pm 122	762 \pm 40	768 \pm 63

575 **D.2 Learning Curves**

576 In Figures 6 through 9, the mean training and test success rates are represented by solid and dashed
 577 lines respectively, with the shaded region indicating the 95% confidence interval. All results are
 578 compiled from 8 random seeds. The data corresponding to the final training steps in these plots
 579 (250M steps for ML-10 and 400M steps for ML-45) are used to present the main results in Table 1.

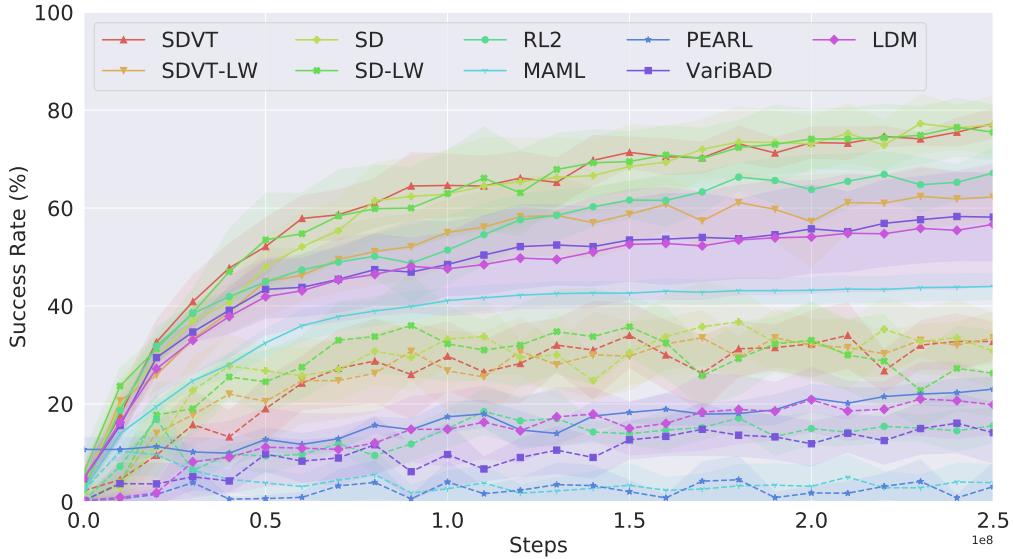


Figure 6: Learning curve on ML-10 – success rate.

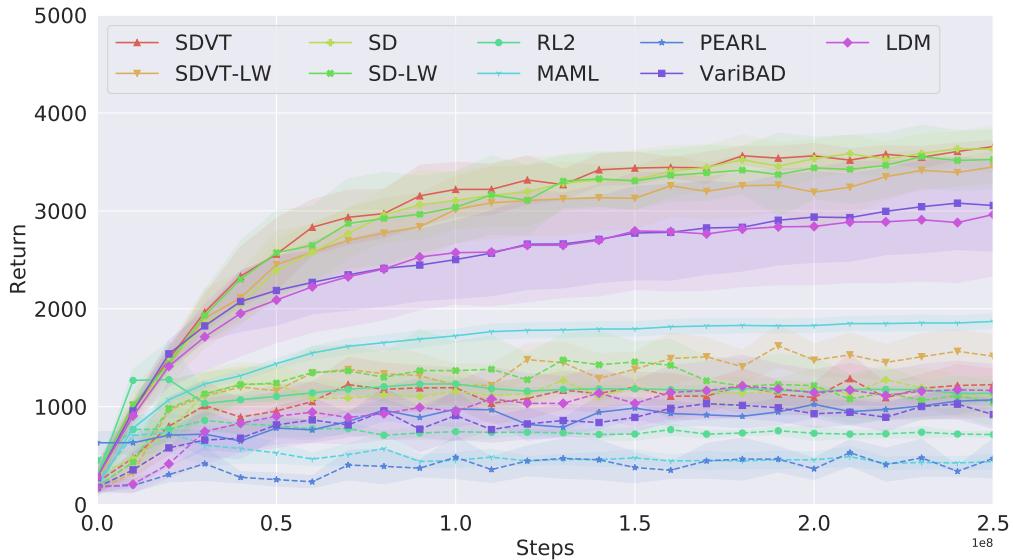


Figure 7: Learning curve on ML-10 – return.

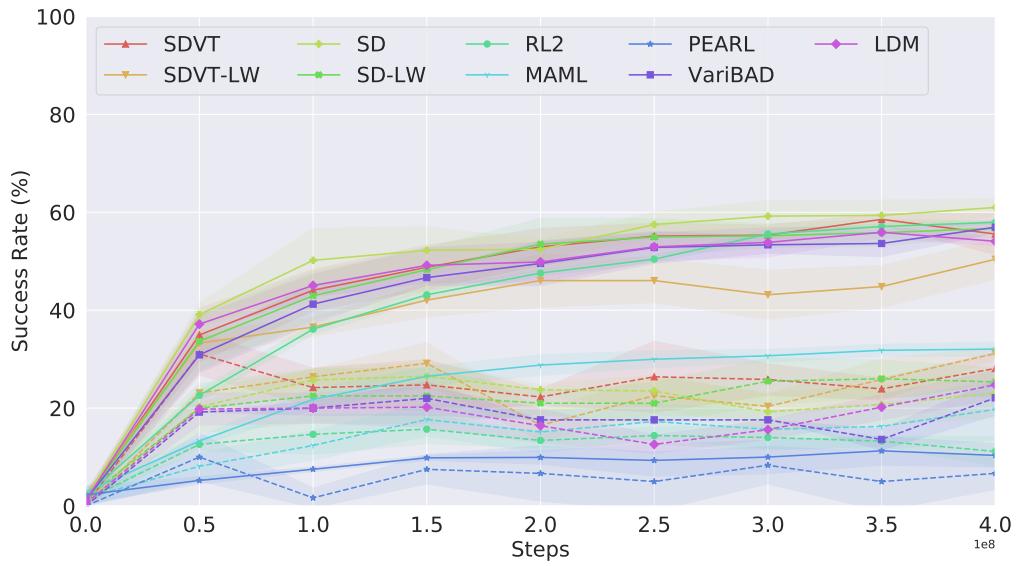


Figure 8: Learning curve on ML-45 – success rate.

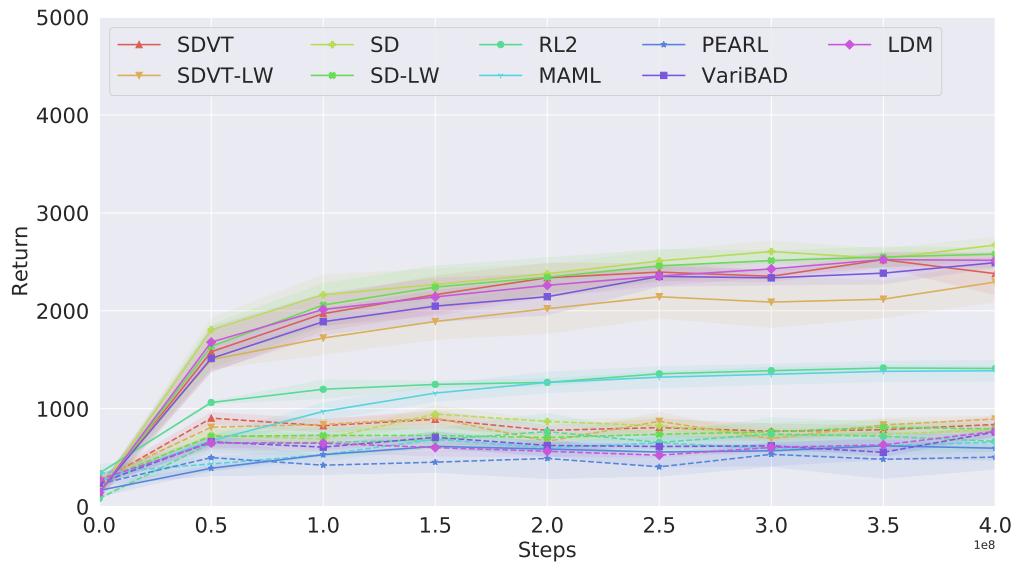


Figure 9: Learning curve on ML-45 – return.

580 **D.3 Subtask Compositions Learned over Training**

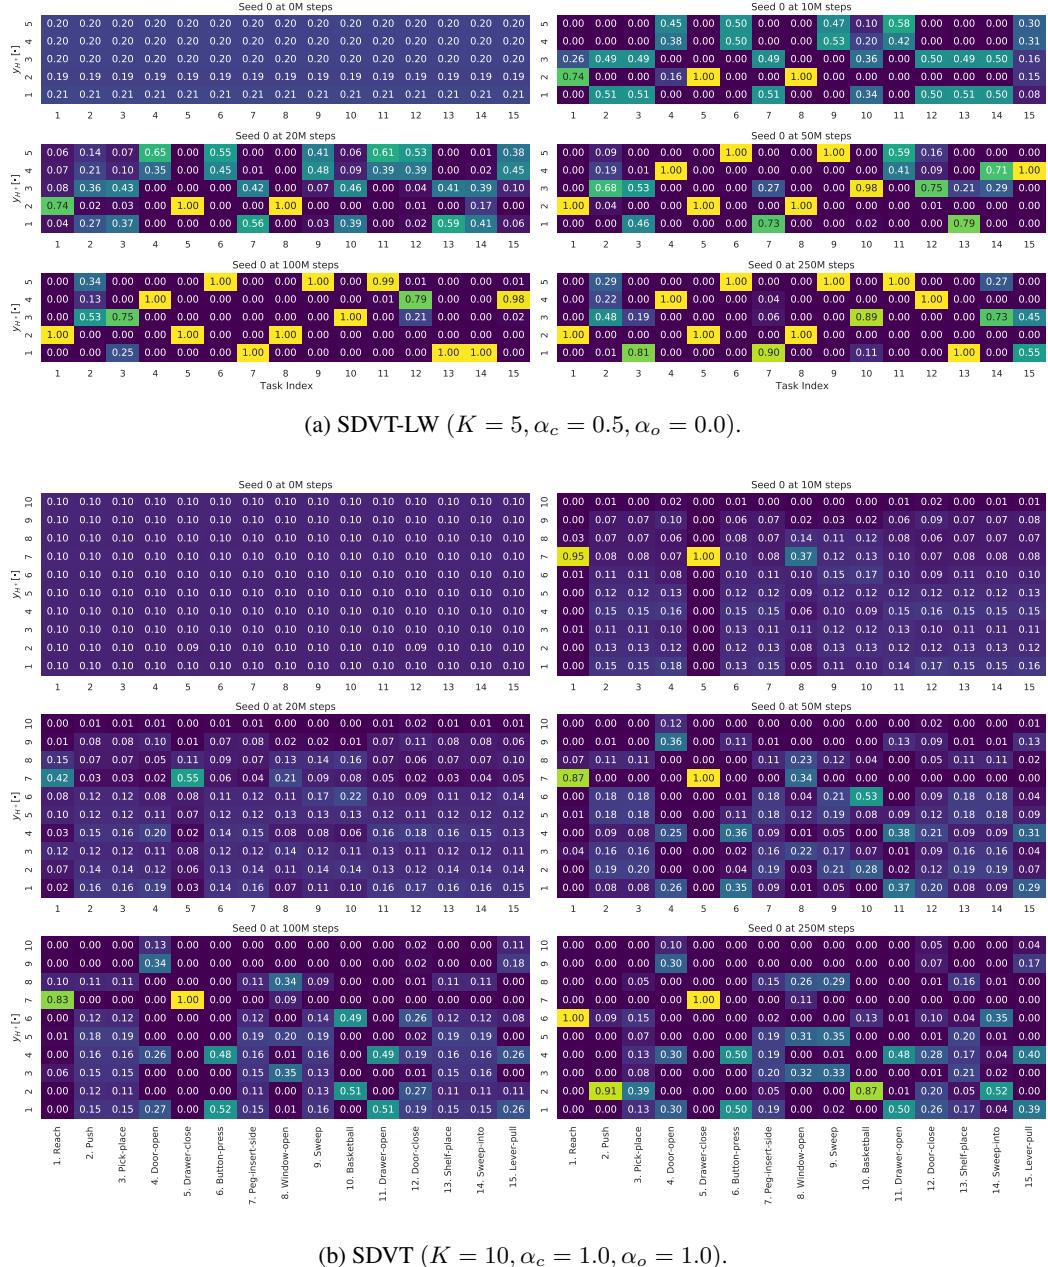


Figure 10: **Subtask compositions learned over training.** We visualize the development of subtask compositions for (a) SDVT-LW and (b) SDVT during training on ML-10. Similar to Figure 4, each column represents the terminal subtask composition (y_{H+}) averaged across 50 parametric variants. (a) As training progresses, subtask compositions of (4) “Door-open” and (12) “Door-close” merge due to shared transition dynamics, which may cause the “Door-close” policy to keep the door open during testing without virtual training, underscoring the crucial role that virtual training can play in these scenarios. (b) With occupancy regularization, as the training progresses, the decomposition process prioritizes occupying lower indices over higher ones in the subtask compositions.

581 D.4 Subtask Compositions of all Seeds

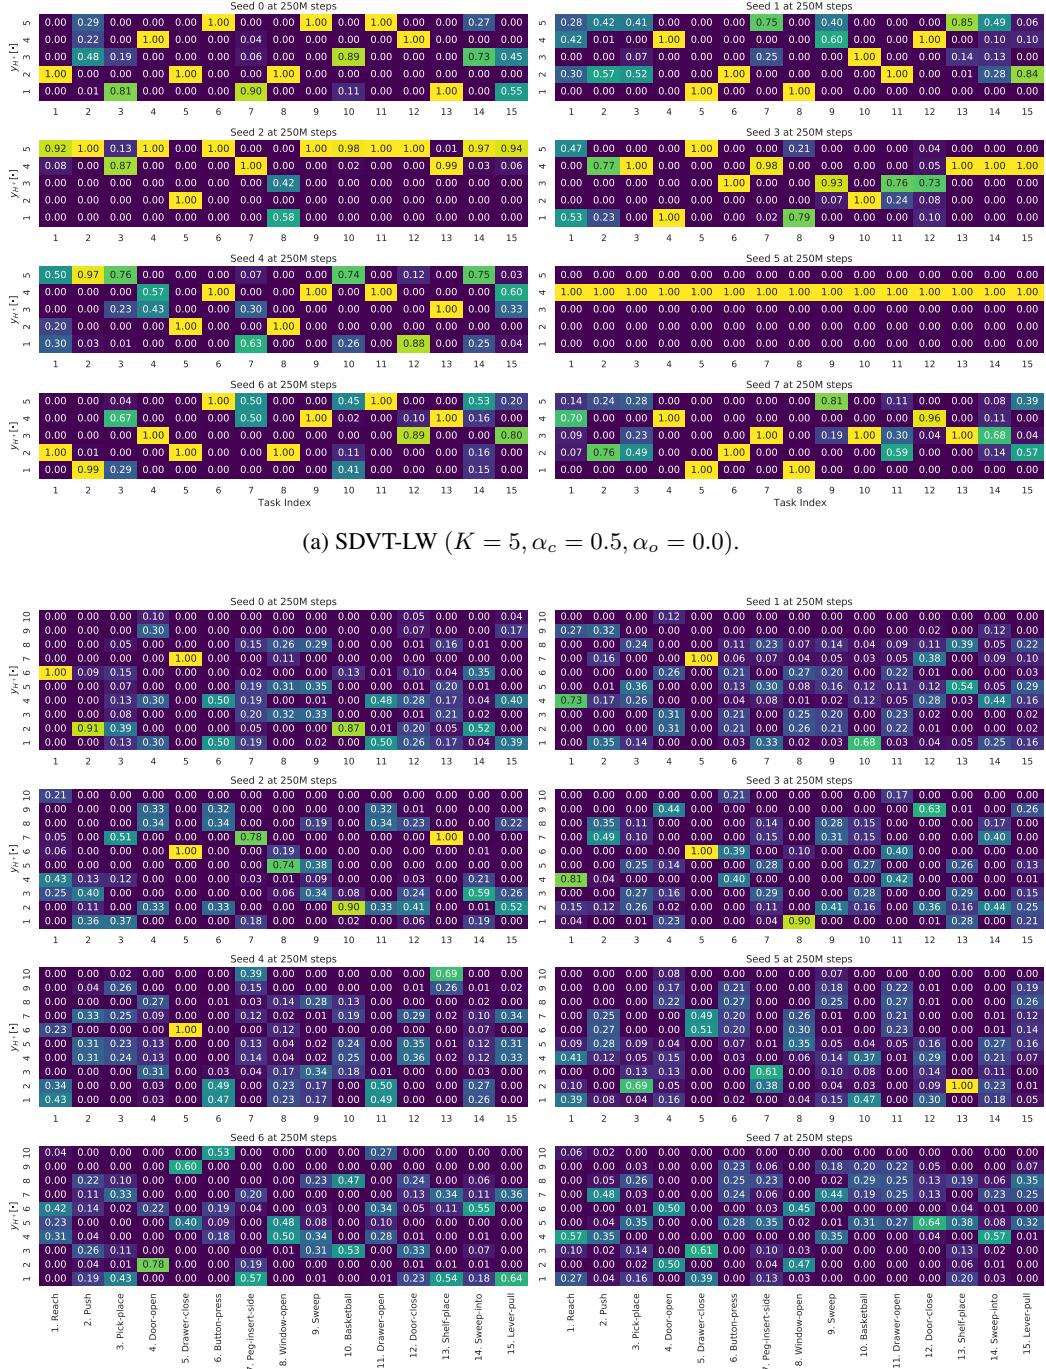


Figure 11: **Subtask compositions of all seeds.** We visualize the subtask compositions of (a) SDVT-LW and (b) on ML-10 after 250M training steps for all seeds. Decomposition processes differ among random seeds due to varying initialization and sample tasks during meta-training. One subtask can be interpreted as a combination of multiple subtasks and vice versa, depending on the seed and the dimension K . We rarely have a collapse to a single Gaussian as in (a) Seed 5, which lowers the average performance. (b) With occupancy regularization, we find that the 10th element of the composition is rarely occupied, whereas the first element is occupied by many tasks.

582 **E Ablation Results**

583 In order to confirm the source of the gain obtained from the proposed method, we conduct an extensive
 584 ablation study based on SDVT and SDVT-LW on ML-10.

585 **E.1 Occupancy Regularization**

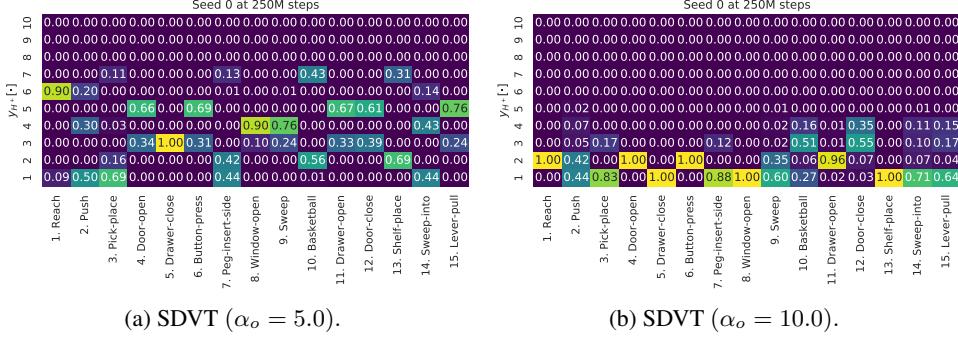


Figure 12: **Occupancy ablation.** We report the learned subtask compositions of SDVT ($K = 10, \alpha_c = 1.0$) for different occupancy coefficients analogous to Figure 4a.

586 Figures 4a and 12 indicate that our occupancy regularization effectively constrains the use of higher
 587 indices as intended. We fix the default choice of $\alpha_o = 1.0$ that we found to work well on ML-10.
 588 Our default SDVT with $\alpha_o = 1.0$ scores a test success rate of 32.8%, which scored the best among
 589 $\alpha_o \in \{0.0, 5.0, 10.0\}$ that scored 22.8%, 25.5%, and 21.0%, respectively.

590 **E.2 Dropout and Dispersion**

Table 10: **Ablation results.** We report the mean success rate (%) and the difference caused by the changes made from the default SDVT-LW ($K = 5, \alpha_c = 0.5, \alpha_o = 0.0$) on ML-10 in parenthesis.

SDVT-LW	without		K			α_c		
	Dropout	Dispersion	3	7	10	0.1	1.0	2.0
ML-10 Train	65.6 (+3.5)	73.0 (+10.9)	60.1 (-2.0)	69.5 (+7.4)	70.8 (+8.7)	59.8 (-2.3)	70.3 (+8.2)	66.3 (+4.2)
	16.5 (-16.9)	21.5 (-11.9)	25.0 (-8.4)	22.0 (-11.4)	21.1 (-12.3)	20.5 (-12.9)	16.1 (-17.3)	17.9 (-15.5)
ML-10 Test	16.5 (-16.9)	21.5 (-11.9)	25.0 (-8.4)	22.0 (-11.4)	21.1 (-12.3)	20.5 (-12.9)	16.1 (-17.3)	17.9 (-15.5)

591 Table 10 demonstrates the significant roles played by both dropout and dispersion in generating
 592 extrapolated dynamics during virtual training, thereby enhancing the test success rate. As outlined in
 593 Appendix E.1 of the LDM paper [33], virtual training without dropout does not exhibit a significant
 594 improvement due to decoder overfitting on the state and action inputs of the training tasks. In fact,
 595 virtual training without dropout can even lead to a decrease in test performance.

596 However, it is important to note that these factors are not the primary contributors to the empirical
 597 gains of our method. Dropout is specifically essential for virtual training, but not necessarily for other
 598 methods that do not employ virtual training, such as VariBAD and RL2 (as discussed in Appendix
 599 E.2 of the LDM paper). We verified this in our ML-10 experiment using VariBAD with dropout. The
 600 inclusion of dropout marginally improves the training success rate of vanilla VariBAD from 58.2%
 601 to 63.0% and the test success rate from 14.1% to 16.5%. However, the performance enhancement
 602 achieved through dropout is not as substantial for VariBAD as it is for SDVT-LW's virtual training.

603 **E.3 Decomposition Distribution**

604 Table 10 emphasizes the significance of hyperparameters such as subtask dimension K and categorical
 605 coefficient α_c in determining subtask compositions. It is crucial to carefully set these hyperparameters

606 based on the number of training tasks and their correlations. In our lightweight (LW) approach, we
 607 have selected $K = 5$ and $\alpha_c = 0.5$, which effectively distributes subtasks across the ML-10 tasks.
 608 When α_c is set to a small value, task classification collapses into a few subtasks, as depicted in Figure
 609 4b. Conversely, with a large value of α_c , the entropy of the composition is maximized, and all tasks
 610 exhibit a uniform probability distribution, as illustrated in Figure 4d. However, it is noteworthy that
 611 SDVT-LW outperforms VariBAD and LDM in test success rate, regardless of the specific values
 612 chosen for K , such as $K = 3, 7, 10$. To alleviate the burden of hyperparameter tuning in our LW
 613 methods, we introduce the occupancy regularization.

614 E.4 Number of parameters

Table 11: **Number of parameters and success rates.** We report the number of parameters used by our methods and baselines. We demonstrate that our gain is not from the increased capacity.

Methods	Number of Parameters				ML-10 Success Rate (%)	
	Encoder	Decoder	Policy	Sum	Train	Test
SDVT-LW	1,047,455	174,821	235,401	1,457,677	62.1	33.4
SD-LW	1,047,455	25,637	235,401	1,308,493	75.5	26.2
LDM	502,580	25,577	202,249	730,406	56.7	19.8
VariBAD	251,290	25,577	202,249	479,116	58.2	14.1
VariBAD (hidden $\times 2$)	894,362	25,577	663,305	1,583,244	62.4	12.0

615 Table 11 indicates that SDVT-LW employs 3 and 2 times more parameters compared to VariBAD
 616 and LDM, respectively. The encoder of our method has more parameters than VariBAD due to the
 617 added categorical layer. The decoder of SDVT-LW includes the parameters of the dispersion layers.
 618 However, we find that the model’s improvement is not mainly attributed to the increased capacity.
 619 Generally, a larger capacity does not necessarily guarantee better generalization. While selecting
 620 encoder (GRU) and policy hidden sizes, we experimented with 128, 256, and 512, discovering that
 621 256 works best for all task-inference methods (LDM, SD only, and SDVT). Notably, VariBAD with
 622 hidden dimensions of 512 possesses more parameters than SDVT but exhibits an even lower test
 623 success rate than the vanilla VariBAD.

624 E.5 Conditioning policy on belief

Table 12: **Masking contexts.** We present the success rates (%) on ML-10, achieved by ablating the contexts fed into the policy, to illustrate the effective utilization of learned contexts by the policy.

	SDVT-LW	SDVT-LW masked (ω_{ϕ_y})	SDVT-LW masked ($\mu_{\phi_z}, \sigma_{\phi_z}$)	SDVT-LW masked ($\omega_{\phi_y}, \mu_{\phi_z}, \sigma_{\phi_z}$)
ML-10 Train	62.1	58.6	36.9	34.8
ML-10 Test	33.4	33.6	25.8	24.5

625 We employ VariBAD’s stop gradient architecture to avoid the instability that may arise from the
 626 concurrent training of policy and ELBO objectives. This may lead to a potential vulnerability where
 627 the policy neglects contexts. However, in Meta-World, identical observations could belong to different
 628 tasks, such as “Reach” and “Pick-place.” Therefore, the agent must rely on the inferred context. To
 629 confirm this, we evaluate SDVT by masking either (ω_{ϕ_y}) or $(\mu_{\phi_z}, \sigma_{\phi_z})$ in the policy input.

630 Our findings reveal that masking the contexts severely damages the training and test success rates.
 631 Especially, masking the Gaussian parameters of z is more critical than masking the categorical
 632 parameter of y . However, this does not suggest that y is redundant, given that the continuous context
 633 z is trained to include y ’s information. Interestingly, the test success rate, with all contexts masked
 634 (i.e., solely relying on the current state s_t), achieves a success rate of 24.5%, surpassing all baseline
 635 scores. Although the policy does not directly employ the context learned by the GMVAE, it can still
 636 be effectively trained in the imaginary task generated by the learned GMVAE to prepare for tests with
 637 unseen compositions of subtasks.