

Not All Tasks Are Equally Difficult: Multi-Task Deep Reinforcement Learning with Dynamic Depth Routing

Jinmin He^{1,2}, Kai Li^{1,2*}, Yifan Zang^{1,2}, Haobo Fu⁶, Qiang Fu⁶, Junliang Xing^{5*}, Jian Cheng^{1,3,4}

¹Institute of Automation, Chinese Academy of Sciences

²School of Artificial Intelligence, University of Chinese Academy of Sciences

³School of Future Technology, University of Chinese Academy of Sciences

⁴AiRiA

⁵Tsinghua University

⁶Tencent AI Lab

{hejinmin2021, kai.li, zangyifan2019, jian.cheng}@ia.ac.cn, {haobofu, leonfu}@tencent.com, jlxing@tsinghua.edu.cn

Abstract

Multi-task reinforcement learning endeavors to accomplish a set of different tasks with a single policy. To enhance data efficiency by sharing parameters across multiple tasks, a common practice segments the network into distinct modules and trains a routing network to recombine these modules into task-specific policies. However, existing routing approaches employ a fixed number of modules for all tasks, neglecting that tasks with varying difficulties commonly require varying amounts of knowledge. This work presents a Dynamic Depth Routing (D2R) framework, which learns strategic skipping of certain intermediate modules, thereby flexibly choosing different numbers of modules for each task. Under this framework, we further introduce a ResRouting method to address the issue of disparate routing paths between behavior and target policies during off-policy training. In addition, we design an automatic route-balancing mechanism to encourage continued routing exploration for unmastered tasks without disturbing the routing of mastered ones. We conduct extensive experiments on various robotics manipulation tasks in the Meta-World benchmark, where D2R achieves state-of-the-art performance with significantly improved learning efficiency.

Introduction

Deep Reinforcement Learning (RL) has witnessed remarkable advancements in the past decade, demonstrating successful applications in diverse domains, including game playing (Mnih et al. 2015; Ye et al. 2020) and robotic control (Lillicrap et al. 2016; Levine et al. 2016). However, despite the current capabilities of deep RL methods to learn individual policies for specific tasks, training a single policy that can generalize across multiple tasks remains highly challenging (Teh et al. 2017; Yu et al. 2020b).

Multi-Task Reinforcement Learning (MTRL) aims to acquire an effective policy capable of accomplishing a diverse range of tasks. Given the potential similarities among the joint learning tasks, training a multi-task policy allows for sharing and reusing knowledge across related tasks, resulting in enhanced sample efficiency compared to training each




Figure 1: The easy task *drawer-close* can be solved by simply pushing at any position of the drawer, so D2R learns to use only 6 modules. However, when tackling the difficult task *drawer-open*, which involves first grasping the handle and then pulling outward, D2R adapts to utilize 10 modules.

task separately (Yang et al. 2020). However, even within the same domain, the variability between tasks still exists, making it challenging to determine what knowledge should be shared between tasks and how to share it appropriately.

To tackle this problem, a class of approaches based on modularization and routing has been proposed (Fernando et al. 2017; Yang et al. 2020; Sun et al. 2022). These approaches usually consist of two networks: a base module network, which divides all parameters into separate modules to learn different knowledge, and a routing network, which determines how these modules are combined. These two networks are jointly optimized over multiple tasks.

Although these existing approaches are promising, their routing networks always use the same number of modules for all tasks. It could result in unnecessary parameter optimization because *not all tasks are equally difficult*. Different tasks demand varying amounts of knowledge, with easier tasks typically requiring fewer modules than difficult ones. Based on this consideration, we introduce a Dynamic Depth Routing (D2R) framework that combines a diverse number of modules with flexible adaptation to task difficulty (Fig.1). Specifically, the modules within D2R’s base module network are arranged in a predetermined topological sequence. For each task, each module establishes routing connections with selected preceding modules. Thus a module is skipped in case any of its following modules does not select it. In this way, D2R automatically determines which modules to use and the number of modules required for each task.

*Corresponding authors.

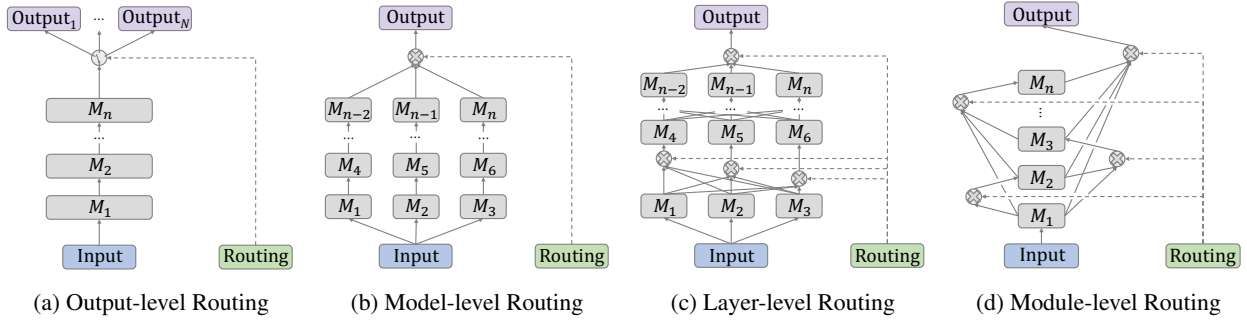


Figure 2: Four different levels of routing with n modules. (a) is the basic multi-head approach. (b) routes from several separate networks. (c) establishes connections between adjacent layers. (d) can build any possible connections to form a DAG routing.

Our D2R framework is agnostic to the underlying RL algorithms. However, during off-policy training, it faces a unique challenge of disparate routing paths between behavior and target policies. To address this issue, we propose a ResRouting method, which avoids potential negative knowledge transfer by prioritizing updating the behavior policy routing without optimizing unsuitable modules for the target policy. In addition, D2R also suffers from the prevalent imbalanced learning problem during multi-task training, wherein easy tasks typically converge faster than difficult ones. To address this problem, we propose an automatic route-balancing mechanism to harmonize the learning process of different tasks by encouraging continued routing exploration for unmastered tasks while simultaneously maintaining the routing stability for already mastered tasks.

We conduct extensive experiments in the Meta-World benchmark (Yu et al. 2020b), consisting of 50 robotic manipulation tasks. D2R achieves significant improvement in both sample efficiency and final performance compared to state-of-the-art MTRL algorithms. We perform rich analyses to demonstrate that D2R learns suitable routing paths based on task difficulty. We also conduct detailed ablation studies to gain insights into the contribution of each component.

Background

Multi-Task Reinforcement Learning

We consider MTRL setting involving N tasks, with each task \mathcal{T} represented by a finite horizon Markov Decision Process (MDP) (Bellman 1966; Puterman 2014). Each MDP can be characterized by $\langle S, A, P, R, \gamma \rangle$, where states $s \in S$ and actions $a \in A$ are continuous in this paper. In addition, $P(s_{t+1}|s_t, a_t)$ denotes the stochastic transition dynamics, $R(s_t, a_t)$ represents the reward function, and γ is the discount factor. In the context of MTRL, the primary objective is to learn a policy $\pi(a_t|s_t)$ that maximizes the expected return across all tasks sampled from a task distribution $p(\mathcal{T})$.

Soft Actor-Critic

In this work, we use the Soft Actor-Critic (SAC) (Haarnoja et al. 2018) algorithm to train a policy across multiple tasks. SAC is an off-policy actor-critic method that leverages the maximum entropy framework. The critic network

$Q_\theta(s_t, a_t)$, representing a soft Q-function, is parameterized by θ , and the actor network $\pi_\phi(a_t|s_t)$ is parameterized by ϕ . The objective of policy optimization is:

$$J_\pi(\phi) = \mathbb{E}_{s_t \sim \mathcal{D}} [\mathbb{E}_{a_t \sim \pi_\phi} [\alpha \log \pi_\phi(a_t|s_t) - Q_\theta(s_t, a_t)]] , \quad (1)$$

where \mathcal{D} represents the data in the replay buffer, and α is a learnable temperature parameter to penalize entropy. It is learned to maintain the entropy level of the policy following:

$$J(\alpha) = \mathbb{E}_{a_t \sim \pi_\phi} [-\alpha \log \pi_\phi(a_t|s_t) - \alpha \bar{\mathcal{H}}] , \quad (2)$$

where $\bar{\mathcal{H}}$ is a desired minimum expected entropy. If the optimization leads to an increase in $\pi_\phi(a_t|s_t)$ with a decrease in the entropy, the temperature α will accordingly increase.

Modularization and Routing

Training a multi-task policy with global parameter sharing poses a challenging optimization problem due to potential conflicts in gradient descent across different tasks (Yu et al. 2020a; Liu et al. 2021). *Modularization* involves dividing parameters into different modules, which reduces gradient interference by utilizing different modules and enables better generalization (Singh 1992; Andreas, Klein, and Levine 2017). We view module reuse as a form of knowledge sharing. Considering that different tasks and states may necessitate varying knowledge, it is natural to require a rational selection and combination of modules, which is referred to as *routing*. To better position our approach, we categorize the existing routing frameworks in Fig.2(a), (b), and (c).

Output-level routing (Fig.2(a)) is the most basic and standard multi-head approach, where tasks share all the modules in the backbone, but each task has its specific head module. It is a fixed routing with no learning process, as each head deterministically matches a specific task.

Model-level routing (Fig.2(b)) involves routing from several separate networks comprising multiple modules, which shares a similar idea with the Mixture-of-Experts (MoE) architecture (Jacobs et al. 1991; Shazeer et al. 2017). In this context, the routing network learns a weight assignment for separate networks, which is analogous to the gating network in MoE used to select a combination of experts.

Layer-level routing (Fig.2(c)) is an extension of model-level routing. It enables each layer to apply separate routing and establish routing connections between modules in two

adjacent layers (Yang et al. 2020), but connections across layers are not allowed. The routing network learns the combining weights of each module within each layer.

Method

The computational cost remains consistent among various tasks in all three routing types mentioned above since the amount of parameter usage remains unchanged across tasks. However, easy tasks only require simpler policies with fewer parameters than difficult ones. Instead of dividing modules into sets and requiring routing between specified sets, we treat each module as an independent unit and propose a new **module-level routing** (Fig.2(d)). In this paradigm, a predetermined topological order organizes all modules. The routing network learns a separate Directed Acyclic Graph (DAG) for each task that satisfies a subsequence of this ordering. In each DAG, the vertices represent the utilized module, and the directed edges represent the learned routing connections between these modules.

In terms of flexibility, output-level routing is the least flexible since its routing path is static and solely determined by task identity. In contrast, the other three routing types allow for more flexible knowledge sharing because their routing paths dynamically change during the training process, and they can incorporate additional clues, such as the current state information. Regarding expressive capability, module-level routing stands out as the strongest. Both model-level and layer-level routing paths can be viewed as specific DAGs with structural restrictions. For instance, layer-level routing prohibits cross-layer connections between non-adjacent layers. Therefore, they are special cases of module-level routing, which does not constrain DAG’s structure.

Dynamic Depth Routing

Considering the varying knowledge requirements for tasks of different complexities, our model aims to select appropriate modules and combine them optimally based on the task difficulty. To achieve this goal, we introduce a Dynamic Depth Routing (D2R) framework, which builds upon module-level routing. As shown in Fig.3, D2R contains two networks: a base module network and a routing network. In each timestep, the routing network recombines the base modules using the current observed state s_t and the task \mathcal{T} . In this work, we use a one-hot vector to represent each task.

Base Module Network. Suppose we have n modules in total, and all modules are arranged in a predetermined topological sequence. The input to the first module is the state representation $F(s_t)$. For the remaining, the input of the i -th module is calculated as a weighted sum of the outputs from previous modules $\{m^j : j < i\}$, determined by a routing probability p^j which is generated by the routing network. The output of the i -th module is represented as:

$$m^i = M^i \left(\sum_{j=1}^{i-1} p_j^i \cdot m^j \right), \quad (3)$$

where M^i is the network of the i -th module. The output dimensions are consistent across all modules except for the

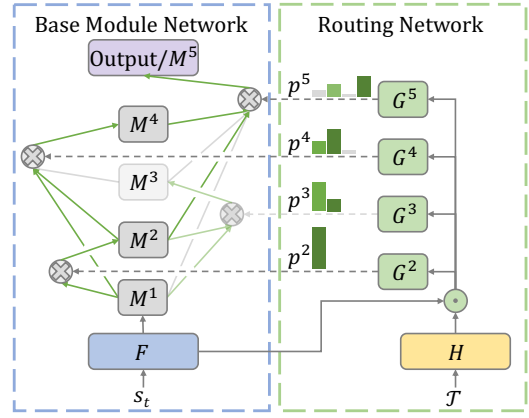


Figure 3: The structure of D2R contains a base module network (left) with multiple modules and a routing network (right) that generates the routing probabilities p^i for each module to select and combine its routing sources.

last one m^n , which should align with the action distribution.

Routing Network. The routing network consists of n sub-networks, each generating routing probability p^i for its corresponding base module. The input to the routing network is the element-wise multiplication of the state representation $F(s_t)$ and the task representation $H(\mathcal{T})$. We define the routing logit of module i as $z^i \in \mathbb{R}^{i-1}$, represented as:

$$z^i = G^i (F(s_t) \cdot H(\mathcal{T})), \quad (4)$$

where G^i is the routing sub-network for the i -th module.

Not all tasks are equally difficult, with easy tasks commonly requiring fewer modules. Therefore, not every task necessitates the full connection of all modules. In addition, in order to train the network in an end-to-end manner, the combination process is required to be differentiable. Thus, we use Top-K MaskSoftmax function to normalize routing logit z^i into routing probability p^i as:

$$p_j^i = \text{MaskSoftmax} (z^i, d^i)_j = \frac{\exp(z_j^i) \cdot d_j^i}{\sum_{k=1}^{i-1} \exp(z_k^i) \cdot d_k^i}, \quad (5)$$

where d^i is the routing path of module i represented as:

$$d_j^i = \text{TopK} (z^i, k)_j = \begin{cases} 1 & \text{if } z_j^i \text{ in top } k \text{ of } z^i. \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

The deterministic routing from top k modules always amounts to an exploitative strategy, resulting in different tasks converging to the same routing paths. To promote exploration and discover improved routing paths, we modify Eq.6 by sampling k routing sources without replacement from the multinomial distribution in the training phase as:

$$d_j^i = \text{SampleK} (z^i, k)_j = \begin{cases} 1 & \text{if } z_j^i \text{ in sampled } k. \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

If the number of preceding modules is less than k , the routing sources include all of them. Referring to the setting of the sparse MoE model (Shazeer et al. 2017; Lepikhin et al. 2021), we set $k = 2$ in our implementation.

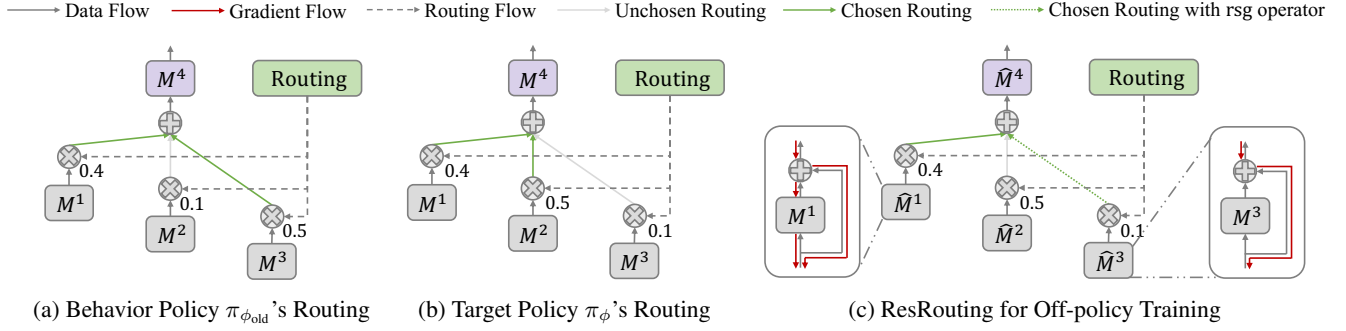


Figure 4: (a) and (b) illustrate the disparate routing paths between behavior policy and target policy during off-policy training. (c) shows the structure of ResRouting, where routing sources with low probabilities are processed using the rsg operator.

As shown in Fig.3, five modules $\{M^i\}_{i=1}^5$ are included, each of which gets its own routing sources indicated by the green arrows. However, M^3 is actually unutilized because the routing path back from output does not go through M^3 (neither M^5 nor M^4 chooses routing from M^3). Hence D2R strategically skips the calculation of certain base modules.

The critic network adopts a similar architecture as the actor network described above, with an additional input parameter (action a_t) and a different output of Q-value.

ResRouting for Off-policy Training

D2R is a generic framework agnostic to the underlying RL algorithms, supporting both on-policy and off-policy approaches. As sample efficiency is a primary motivation for MTRL, off-policy algorithms are more commonly used (Yu et al. 2020b), which learn with substantially fewer environment samples than on-policy ones. However, D2R faces a unique challenge during off-policy training. Specifically, the behavior policy $\pi_{\phi_{old}}$, which generates training trajectories, may follow different routing paths from the target policy π_{ϕ} , which is currently evaluated and improved. Fig.4 illustrates a top-2 routing sample where module M^4 chooses to route from M^1 and M^3 due to the highest routing probabilities in $\pi_{\phi_{old}}$ (Fig.4(a)); for now, M^1 and M^2 are chosen by π_{ϕ} (Fig.4(b)). Thus, a key question is how to update the target policy under this routing path inconsistency.

We propose a ResRouting method (Fig.4(c)) to address this challenge. Given that $\pi_{\phi_{old}}$ generates actions through the routing path d_{old} to interact with the environment, it is natural to evaluate and update the corresponding routing strategy. Therefore, we store the routing paths d_{old} of $\pi_{\phi_{old}}$ in the replay buffer and utilize them to train π_{ϕ} . The routing probability of π_{ϕ} in Eq.5 is modified as:

$$p_j^i = \text{MaskSoftmax}(z^i, d_{old}^i)_j. \quad (8)$$

However, blindly optimizing all the modules in d_{old} may cause negative knowledge transfer since some modules in d_{old} might no longer be suitable in π_{ϕ} since it has been trained for a while. The optimization of these unsuitable modules does not yield performance improvement; on the contrary, it adversely impacts the other paths that utilize these modules. Therefore, we prevent the update of these unsuitable modules in π_{ϕ} . The output of each module in Eq.3

is modified as:

$$m^i = M^i \left(\sum_{j=1}^{i-1} p_j^i \cdot \chi(m^j, z_j^i, \text{sg}) \right), \quad (9)$$

where sg stands for the stop-gradient operator, defined as an identity function during forward computation with zero partial derivatives. $\chi(m^j, z_j^i, \text{sg})$ is an indicator function used to decide the application of the sg operator on module m^j ,

$$\chi(m^j, z_j^i, \text{sg}) = \begin{cases} m^j & \text{if } \text{Softmax}(z_j^i) \geq \sigma^i, \\ \text{sg}[m^j] & \text{if } \text{Softmax}(z_j^i) < \sigma^i, \end{cases} \quad (10)$$

where threshold σ^i relates to the number of optional routing sources, which is set to $1/i$. As shown in Fig.4, M^3 is unsuitable for M^4 when updating π_{ϕ} ($0.1 < 1/4$) even though it was previously selected by $\pi_{\phi_{old}}$.

However, apart from blocking the gradients of M^3 , the original sg operator also prevents updating all M^3 's routing source modules. Therefore, later modules are more likely to be optimized than the preceding ones. It raises a subtle optimization problem of reduced training efficiency and unbalanced updating frequency. Inspired by the residual learning in ResNet (He et al. 2016), ResRouting adopts a similar structure with an additional shortcut connection. For suitable modules (like M^1 in Fig.4(c)), ResRouting preserves the full gradient information. However, for unsuitable modules (like M^3 in Fig.4(c)), ResRouting stops the gradient to avoid updating these specific modules while retaining the gradient information in the shortcut connection, thereby facilitating the update of preceding modules. So Eq.9 is modified as:

$$m^i = \sum_{j=1}^{i-1} p_j^i \cdot \hat{m}^j + M^i \left(\sum_{j=1}^{i-1} p_j^i \cdot \hat{m}^j \right), \quad (11)$$

where $\hat{m}^j = \chi(m^j, z_j^i, \text{rsg})$, and rsg operator is defined as:

$$\text{rsg}[m^i] = \sum_{j=1}^{i-1} p_j^i \cdot \hat{m}^j + \text{sg} \left[M^i \left(\sum_{j=1}^{i-1} p_j^i \cdot \hat{m}^j \right) \right]. \quad (12)$$

In summary, ResRouting satisfactorily prioritizes updating the routing of the behavior policy while avoiding updating unsuitable modules in the current target policy simultaneously. Furthermore, it enables the effective training of all modules, thereby enhancing training efficiency.

Automatic Route-Balancing

Besides the off-policy routing inconsistency issue, D2R faces another challenge of an imbalanced learning process when simultaneously optimizing multiple tasks, where easy tasks typically converge faster than difficult ones. To this end, we propose an automatic route-balancing mechanism to harmonize the learning process of different tasks. It encourages continued exploration of routing paths for the unmastered difficult tasks while exploiting routing paths for the mastered easy ones. Specifically, we introduce a hyper-parameter $\tau_{\mathcal{T}}$ to the routing path sampling for each task \mathcal{T} ,

$$d_j^i = \text{SampleK}(z^i/\tau_{\mathcal{T}}, k)_j, \quad (13)$$

where $z^i/\tau_{\mathcal{T}}$ replaces z^i in Eq.7. For unmastered tasks, $\tau_{\mathcal{T}}$ should be larger to make the routing distribution smoother, thus facilitating exploration. In contrast, for mastered tasks, $\tau_{\mathcal{T}}$ should be smaller to emphasize exploitation.

Manually specifying $\tau_{\mathcal{T}}$ for each task is cumbersome and time-consuming. Inspired by (Yang et al. 2020), we can automatically adjust $\tau_{\mathcal{T}}$ online by leveraging the favorable properties of the SAC temperature parameter α in Eq.2 as:

$$\tau_{\mathcal{T}} = \frac{1/\alpha_{\mathcal{T}}}{\sum_{j=1}^N (1/\alpha_j)}, \quad (14)$$

where $\alpha_{\mathcal{T}}$ is the SAC temperature parameters for task \mathcal{T} , and N is the total number of tasks. For unmastered (mastered) tasks, their policy entropy tends to be high (small). SAC decreases (increases) the value of α through Eq.2, which, in turn, increases (decreases) the value of $\tau_{\mathcal{T}}$ according to Eq.14, meeting our requirements. The full pseudocode for D2R is described in Appendix B (Algorithm 1).

Related Work

Multi-task Reinforcement Learning. Multi-task learning is a training paradigm that improves generalization by utilizing the domain information contained in potentially related tasks (Caruana 1997). MTRL extends the idea of multi-task learning to reinforcement learning, expecting that knowledge shared across tasks will be discovered by simultaneously learning multiple RL tasks. It has been extensively investigated from various perspectives (Calandriello, Lazaric, and Restelli 2014; Hessel et al. 2019; Xu et al. 2020). (Yu et al. 2020a; Liu et al. 2021) view MTRL as a multi-objective optimization problem aiming to handle conflicting gradients arising from different task losses during training, but they lead to an increase in computational and memory consumption. (Parisotto, Ba, and Salakhutdinov 2016; Teh et al. 2017) utilize policy distillation to integrate separate policies for different tasks into a unified policy, yet they still require separate networks for different policies and an additional distillation phase. (Sodhani, Zhang, and Pineau 2021; Cho, Jung, and Sung 2022) attempt to choose or learn better representations as more effective task-condition for policy training. In addition, (Sodhani, Zhang, and Pineau 2021) can be viewed as a model-level routing approach because it leverages a mixture of encoders to learn a context-dependent representation effectively. (Yang et al.

2020) stands for layer-level routing, which generates the combination weight for modules between layers. Instead of routing for the outputs of modules, (Sun et al. 2022) explores a compositional parameter space to reconfigure the parameters of each task’s network. All of the above routing approaches limit the number of parameters per task to remain the same. Differently, D2R proposes a more flexible routing scheme that can reasonably allocate the number of modules among different tasks according to the task difficulty.

Dynamic Neural Network. The classical dynamic network architecture Mixture-of-Experts (MoE) (Jacobs et al. 1991; Ma et al. 2018) builds multiple networks as experts in parallel, performing inference with dynamic width. Sparse MoE (Shazeer et al. 2017; Lepikhin et al. 2021; Fedus, Zoph, and Shazeer 2022) inspires our work that the amount of computation can remain relatively small compared to the total model size. Our work is more related to studies on dynamic neural networks with dynamic depth. The early-exiting paradigm allows easy samples to be output at shallow exits without executing deeper layers (Teerapittayanon, McDanel, and Kung 2016; Bolukbasi et al. 2017), but the cross-layer connections are not permitted. Instead of skipping the execution of all the deep layers after a certain classifier, methods like (Wang et al. 2018; Veit and Belongie 2018) selectively skip intermediate layers conditioned on samples. Besides optionally skipping some modules, D2R also dynamically recombines routing sources for each module.

Experiments

Experimental Setup

Benchmarks. We evaluate D2R on Meta-World (Yu et al. 2020b), an MTRL benchmark consisting of 50 robotics manipulation tasks with a sawyer arm in the MuJoCo environment (Todorov, Erez, and Tassa 2012). It offers two setups: **MT10** (a suite of 10 tasks) and **MT50** (a suite of 50 tasks). Following (Yang et al. 2020), we extend all tasks to a random-goal setting. Specifically, **MT10-Fixed** and **MT50-Fixed** refer to tasks with fixed goals, while **MT10-Rand** and **MT50-Rand** denote tasks with random goals.

Baselines. We compare D2R against five baselines. (i) **Multi-task SAC:** Extension of SAC for MTRL with one-hot encoding for task representation. (ii) **Multi-head SAC:** SAC with a shared network backbone apart from independent heads for each task. (iii) **SM** (Yang et al. 2020): Learning a routing network for soft combination of modules. (iv) **CARE** (Sodhani, Zhang, and Pineau 2021): Leveraging additional metadata and a mixture of encoders for task representation. (v) **PaCo** (Sun et al. 2022): A parameter compositional approach to recombine task-specific parameters.

Training Settings. We train all methods with 15 million environment steps (1.5 million per task) using 10 parallel environments on the MT10 setting and 50 million environment steps (1 million per task) using 50 parallel environments on the MT50 setting. Because effective knowledge sharing between tasks improves sampling efficiency and thus accelerates convergence, we reduce the number of sampling environment steps per task as the number of tasks increases.




Figure 5: Comparison of D2R against baselines on four benchmark settings.

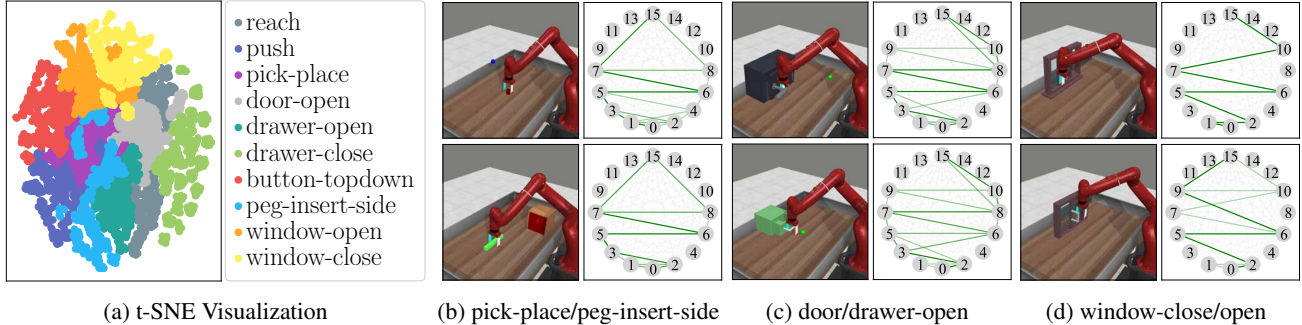


Figure 6: (a): Visualization of 1,000 samples per task with t-SNE. (b)-(d): Comparing routing paths of D2R with complete 16 modules. The light gray fully connected graph displays all possible connections, while the green edges represent the routing paths, where darker color means higher routing probability. Each column illustrates two different tasks.

Quantitative Result

The evaluation metrics of the learned universal policy are based on the success rate of task execution, which is well-defined in the Meta-World benchmark. We evaluate each method every 10K steps per task with 32 episodes and report the 1st, median, and 3rd quartile success rates across 5 different seeds, as illustrated in Fig.5. D2R is closely related to SM, a typical layer-level routing method. SM consists of 4 layers, each containing 4 modules. Thus, we use the same setup of 16 modules to ensure a fair comparison. In both MT10 settings, D2R exhibits impressive performance and stands out as the only method achieving nonzero successful rates across all ten tasks. We also observe that in MT50 settings, D2R demonstrates improved success rates and faster convergence compared to the baseline methods. For the final performance evaluation, we assess the final policy through 100 episodes per task and present the results alongside the parameter size of each method in Table.8 (Appendix C.1).

Routing Analysis

To gain a deeper understanding of D2R’s behavior, we analyze its routing on MT10-Rand from three aspects.

t-SNE Visualization. We evaluate the final policy on each task five times, collecting 1,000 samples of routing paths per task. We reduce the dimension and visualize the routing paths with t-SNE (Van der Maaten and Hinton 2008) in Fig.6(a). The result demonstrates clear boundaries between tasks, showcasing the agent’s ability to distinguish different

tasks. Additionally, we observe intersections among certain tasks as they share common skill sets. For instance, *[push, pick-place, peg-insert-side]* and *[window-open, window-close]* exhibit overlaps. Furthermore, certain tasks, such as *reach*, are divided into multiple sets because they represent fundamental skills that most modules need to handle.

Routing Path Visualization. We present a comparison of complete routing paths for different tasks in Fig.6(b), (c), and (d). Each image pair illustrates the state at a specific timestep on the left side and the corresponding routing path at that moment on the right. The comparison between task *pick-place* and *peg-insert-side* (Fig.6(b)) shows similar routing paths, as they both involve the skill of moving an object from one location to another, with the exception that *pick-place* utilizes an additional module M^4 . In another comparison, task *door-open* and *drawer-open* (Fig.6(c)) share nearly identical routing paths during the open operation. Likewise, task *window-close* and *window-open* (Fig.6(d)) exhibit similar routing paths. Interestingly, we find that the routing path of *window-close* eventually converges to single routing connections, which does not require inter-module information fusion despite using the Top-2 routing function.

Routing in Relation to Task Difficulty. We count the number of modules utilized in 1,000 samples per task and report the mean count for each task with standard derivations in Fig.8. It is evident that difficult tasks (e.g., *peg-insert-side* uses 10.6 modules on average) indeed use more modules than easy ones (e.g., *drawer-close* uses 6.0 on average).




Figure 7: Four distinct ablation studies on the MT10-Rand setting.




Figure 8: Number of modules utilized in each task.

Ablation Study

Effect of Top-K Routing Function. Besides the Top-K Routing function we employ, D2R is also compatible with other routing functions. We show the performance of D2R under various routing functions in Fig.7(a). Hard Routing, also known as Top-1 Routing, selects a single routing source for each module. To address the non-differentiable problem of single path routing, we retain the routing probability of the Softmax function for all routing sources without rescaling the selected one (Fedus, Zoph, and Shazeer 2022). In contrast, Soft Routing establishes all potential routing connections to form a maximum DAG routing path. Hard Routing restricts modules to single connections without module knowledge fusion information, while Soft Routing wastes computing resources on optimizing unnecessary modules. Consequently, both routing functions perform worse than our Top-2 approach. We find that Soft Routing also converges to a sparse path for each input. The detailed discussion on this finding in Appendix C.3 demonstrates that Top-2 Routing is beneficial for enhancing learning efficiency by pruning insignificant dense routing during training.

Effect of SampleK Routing and ResRouting Method. In Fig.7(b), we investigate the impact of different routing strategies on D2R during off-policy training. When we exclude ResRouting (D2R *w/o* ResRouting), there is a significant decrease in the success rate and an increase in variance. Hence, ResRouting plays an important role in mitigating negative knowledge transfer when training D2R with the behavior policy routing paths. Since the target policy can learn from data generated by a completely unrelated behavior policy (*e.g.*, from a human expert) (Sutton and Barto 2018), it can adopt its own routing paths instead of emulating the behavior policy. Thus, we also test D2R *w/* Target Routing, which uses newly sampled routing paths. Its performance and stability are improved compared to D2R *w/o* ResRouting. However, it falls short of D2R because the improved routing paths explored by the behavior policy are not

well-exploited in training. Another variant, D2R *w/o* SampleK, uses TopK (Eq.6) instead of SampleK (Eq.7) during the training phase, but it proves to be detrimental to performance because it lacks exploration of various routing paths.

Effect of Automatic Route-Balancing Mechanism. We present the performance of D2R without the automatic route-balancing mechanism in Fig.7(c). D2R *w/o* Route-Balancing exhibits comparable learning efficiency with D2R in the initial training phase. However, as the agent becomes proficient in certain tasks, D2R with the route-balancing mechanism encourages enhanced exploration of routing paths for the tasks that are not yet mastered, thereby leading to improved overall performance.

Effect of Module-Level Routing. We reimplement both model-level and layer-level routing types by integrating all key components of D2R, *i.e.*, the Top-K routing function, the SampleK routing exploration, the ResRouting method, and the automatic route-balancing mechanism. The results, depicted in Fig.7(d), demonstrate that as the routing flexibility and the expressive capability improve, the corresponding method also shows improved performance. D2R’s module-level routing offers the highest flexibility and the strongest expressive capability, leading to the best performance.

Conclusion and Discussion

In this work, we start by introducing module-level routing, which enables more flexible knowledge sharing. We then propose D2R, designed to adapt varying quantities of knowledge for tasks with differing complexity. To address the challenge of disparate routing paths during off-policy training, we present the ResRouting method as a solution. Furthermore, we design an automatic route-balancing mechanism to balance routing exploration and exploitation for unmastered and mastered tasks. As a result, D2R significantly improves knowledge sharing and outperforms other baseline methods.

One limitation of our proposed approach is that the number of base modules is predetermined, potentially leading to situations where specific modules remain unused by any tasks. Consequently, exploring methods for dynamically generating modules as required by the task presents an intriguing avenue for future research. Furthermore, since modularization provides a natural incremental modeling structure, D2R’s ability for other challenging tasks like few-shot transfer learning will also need investigations.

Acknowledgements

This work is supported in part by the National Key Research and Development Program of China under Grant 2022ZD0116401; the Natural Science Foundation of China under Grant 62076238, Grant 62222606, and Grant 61902402; the Jiangsu Key Research and Development Plan (No. BE2023016); and the China Computer Federation (CCF)-Tencent Open Fund.

References

- Andreas, J.; Klein, D.; and Levine, S. 2017. Modular Multi-task Reinforcement Learning with Policy Sketches. In *Proceedings of International Conference on Machine Learning*, 166–175.
- Bellman, R. 1966. Dynamic Programming. *Science*, 153(3731): 34–37.
- Bolukbasi, T.; Wang, J.; Dekel, O.; and Saligrama, V. 2017. Adaptive Neural Networks for Efficient Inference. In *Proceedings of International Conference on Machine Learning*, 527–536.
- Calandriello, D.; Lazaric, A.; and Restelli, M. 2014. Sparse Multi-Task Reinforcement Learning. In *Advances in Neural Information Processing Systems*, 819–827.
- Caruana, R. 1997. Multitask Learning. *Machine Learning*, 28: 41–75.
- Cho, M.; Jung, W.; and Sung, Y. 2022. Multi-Task Reinforcement Learning with Task Representation Method. In *ICLR Workshop on Generalizable Policy Learning in Physical World*, 1–11.
- Fedus, W.; Zoph, B.; and Shazeer, N. 2022. Switch Transformers: Scaling to Trillion Parameter Models with Simple and Efficient Sparsity. *Journal of Machine Learning Research*, 23(1): 5232–5270.
- Fernando, C.; Banarse, D.; Blundell, C.; Zwols, Y.; Ha, D.; Rusu, A. A.; Pritzel, A.; and Wierstra, D. 2017. PathNet: Evolution Channels Gradient Descent in Super Neural Networks. arXiv:1701.08734.
- Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft Actor-Critic: Off-policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. In *Proceedings of International Conference on Machine Learning*, 1856–1865.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 770–778.
- Hessel, M.; Soyer, H.; Espenholt, L.; Czarnecki, W.; Schmitt, S.; and Van Hasselt, H. 2019. Multi-Task Deep Reinforcement Learning with Popart. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 3796–3803.
- Jacobs, R. A.; Jordan, M. I.; Nowlan, S. J.; and Hinton, G. E. 1991. Adaptive Mixtures of Local Experts. *Neural Computation*, 3(1): 79–87.
- Lepikhin, D.; Lee, H.; Xu, Y.; Chen, D.; Firat, O.; Huang, Y.; Krikun, M.; Shazeer, N.; and Chen, Z. 2021. GShard: Scaling Giant Models with Conditional Computation and Automatic Sharding. In *Proceedings of International Conference on Learning Representations*, 1–14.
- Levine, S.; Finn, C.; Darrell, T.; and Abbeel, P. 2016. End-to-end Training of Deep Visuomotor Policies. *The Journal of Machine Learning Research*, 17(1): 1334–1373.
- Lillicrap, T. P.; Hunt, J. J.; Pritzel, A.; Heess, N.; Erez, T.; Tassa, Y.; Silver, D.; and Wierstra, D. 2016. Continuous Control with Deep Reinforcement Learning. In *Proceedings of International Conference on Learning Representations*, 1–10.
- Liu, B.; Liu, X.; Jin, X.; Stone, P.; and Liu, Q. 2021. Conflict-Averse Gradient Descent for Multi-Task Learning. In *Advances in Neural Information Processing Systems*, 18878–18890.
- Ma, J.; Zhao, Z.; Yi, X.; Chen, J.; Hong, L.; and Chi, E. H. 2018. Modeling Task Relationships in Multi-Task Learning with Multi-Gate Mixture-of-Experts. In *Proceedings of ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 1930–1939*.
- Mnih, V.; Kavukcuoglu, K.; Silver, D.; Rusu, A. A.; Veness, J.; Bellemare, M. G.; Graves, A.; Hiedmiller, M. A.; Fidjeland, A.; Ostrovski, G.; Petersen, S.; Beattie, C.; Sadik, A.; Antonoglou, I.; King, H.; Kumaran, D.; Wierstra, D.; Legg, S.; and Hassabis, D. 2015. Human-Level Control through Deep Reinforcement Learning. *Nature*, 518(7540): 529–533.
- Parisotto, E.; Ba, L. J.; and Salakhutdinov, R. 2016. Actor-Mimic: Deep Multitask and Transfer Reinforcement Learning. In *Proceedings of International Conference on Learning Representations*, 1–9.
- Puterman, M. L. 2014. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q. V.; Hinton, G. E.; and Dean, J. 2017. Outrageously Large Neural Networks: The Sparsely-Gated Mixture-of-Experts Layer. In *Proceedings of International Conference on Learning Representations*, 1–12.
- Singh, S. P. 1992. Transfer of Learning by Composing Solutions of Elemental Sequential Tasks. *Machine Learning*, 8: 323–339.
- Sodhani, S.; Zhang, A.; and Pineau, J. 2021. Multi-Task Reinforcement Learning with Context-based Representations. In *Proceedings of International Conference on Machine Learning*, 9767–9779.
- Sun, L.; Zhang, H.; Xu, W.; and Tomizuka, M. 2022. PaCo: Parameter-Compositional Multi-Task Reinforcement Learning. In *Advances in Neural Information Processing Systems*, 21495–21507.
- Sutton, R. S.; and Barto, A. G. 2018. *Reinforcement Learning: An Introduction*. MIT press.
- Teerapittayanon, S.; McDanel, B.; and Kung, H. 2016. BranchyNet: Fast Inference via Early Exiting from Deep

- Neural Networks. In *Proceedings of International Conference on Pattern Recognition*, 2464–2469.
- Teh, Y.; Bapst, V.; Czarnecki, W. M.; Quan, J.; Kirkpatrick, J.; Hadsell, R.; Heess, N.; and Pascanu, R. 2017. Distral: Robust Multitask Reinforcement Learning. In *Advances in Neural Information Processing Systems*, 4496–4506.
- Todorov, E.; Erez, T.; and Tassa, Y. 2012. Mujoco: A Physics Engine for Model-based Control. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*, 5026–5033.
- Van der Maaten, L.; and Hinton, G. 2008. Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9(86): 2579–2605.
- Veit, A.; and Belongie, S. 2018. Convolutional Networks with Adaptive Inference Graphs. In *Proceedings of the European Conference on Computer Vision*, 3–18.
- Wang, X.; Yu, F.; Dou, Z.-Y.; Darrell, T.; and Gonzalez, J. E. 2018. SkipNet: Learning Dynamic Routing in Convolutional Networks. In *Proceedings of the European Conference on Computer Vision*, 409–424.
- Xu, Z.; Wu, K.; Che, Z.; Tang, J.; and Ye, J. 2020. Knowledge Transfer in Multi-Task Deep Reinforcement Learning for Continuous Control. In *Advances in Neural Information Processing Systems*, 15146–15155.
- Yang, R.; Xu, H.; Wu, Y.; and Wang, X. 2020. Multi-Task Reinforcement Learning with Soft Modularization. In *Advances in Neural Information Processing Systems*, 4767–4777.
- Ye, D.; Liu, Z.; Sun, M.; Shi, B.; Zhao, P.; Wu, H.; Yu, H.; Yang, S.; Wu, X.; Guo, Q.; Chen, Q.; Yin, Y.; Zhang, H.; Shi, T.; Wang, L.; Fu, Q.; Yang, W.; and Huang, L. 2020. Mastering Complex Control in MOBA Games with Deep Reinforcement Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 6672–6679.
- Yu, T.; Kumar, S.; Gupta, A.; Levine, S.; Hausman, K.; and Finn, C. 2020a. Gradient Surgery for Multi-Task Learning. In *Advances in Neural Information Processing Systems*, 5824–5836.
- Yu, T.; Quillen, D.; He, Z.; Julian, R.; Hausman, K.; Finn, C.; and Levine, S. 2020b. Meta-World: A Benchmark and Evaluation for Multi-Task and Meta Reinforcement Learning. In *Proceedings of the Conference on Robot Learning*, 1094–1100.