

Model Adaptation for Time Constrained Embodied Control

Jaehyun Song*, Minjong Yoo*, Honguk Woo[†]

Department of Computer Science and Engineering, Sungkyunkwan University

{s7159540, mjyoo2, hwoo}@skku.edu

Abstract

When adopting a deep learning model for embodied agents, it is required that the model structure be optimized for specific tasks and operational conditions. Such optimization can be static such as model compression or dynamic such as adaptive inference. Yet, these techniques have not been fully investigated for embodied control systems subject to time constraints, which necessitate sequential decision-making for multiple tasks, each with distinct inference latency limitations. In this paper, we present MoDeC, a time constraint-aware embodied control framework using the modular model adaptation. We formulate model adaptation to varying operational conditions on resource and time restrictions as dynamic routing on a modular network, incorporating these conditions as part of multi-task objectives. Our evaluation across several vision-based embodied environments demonstrates the robustness of MoDeC, showing that it outperforms other model adaptation methods in both performance and adherence to time constraints in robotic manipulation and autonomous driving applications.

1. Introduction

In the literature on embodied artificial intelligence (embodied AI), where deep learning models have been increasingly adopted, optimizing the deep learning model structure for specific tasks and operational conditions becomes crucial. Several studies focus on static model optimization and compression [12, 21, 34]. On the other hand, there are a few studies that investigate dynamic model structures, which are designed to adapt more effectively to conditions that change over time [1, 2, 18, 23, 39]. These dynamic approaches, often known as adaptive inference, are particularly beneficial for embodied environments where surrounding conditions evolve and the agent model’s capabilities to adapt in real-time are essential. Yet, a substantial gap remains in the exploration of these approaches within the context of

embodied control systems, especially those operating under strict time constraints and involving more than a single task. Specifically, the systems are distinct in their requirement for sequential decision-making across multiple tasks, each with its own latency limitations. For example, consider an autonomous driving agent. This agent must continuously make driving control decisions in response to moving obstacles. Each decision must be made rapidly and accurately, as any delay can significantly impact driving safety and efficiency due to the time-sensitive nature of driving tasks.

To address the challenges faced in such time-constrained embodied environments, we present MoDeC (MoDeL adaptation for time constrained Embodied Control), a novel modular network framework with efficient and adaptive inference capabilities. In embodied AI, where models are deployed on specific hardware with different limitations, the ability to rapidly adapt to both operational resource conditions and time constraints is crucial. Recognizing this, we employ a constraint-aware modular model architecture, which can transform the procedure of time-sensitive inference into effective module selection within a single modular network that can be deployed on different target devices. This approach enables dynamic model adaptation to varied operational conditions in a sample-efficient way. We also use a meta-learning scheme combined with knowledge distillation for restructuring the module selection in non-iterative time-sensitive computations.

Through intensive experiments with several embodied control scenarios such as robot manipulation tasks in Meta-world [42], autonomous driving tasks in CARLA [7], and object navigation tasks in AI2THOR [16], we demonstrate that our MoDeC framework is applicable for different time constraints and devices, achieving robust adaptation performance in terms of both constraint satisfaction and model accuracy. For instance, MoDeC shows a performance gain of 14.4% in success rates over the most competitive baseline, DS-Net [18] for autonomous driving tasks in CARLA, while it keeps the violation of time constraints to be less than 1%. The main contributions of this paper are summarized as follows.

- We present the constraint-aware modular model frame-

*Equal contribution.

[†]Honguk Woo is the corresponding author.

work MoDeC, specifically designed for dynamic multi-task model adaptation to time constraints and device resource specifications.

- We devise an efficient joint learning algorithm for optimizing the combinatorial module selection in a model and stabilizing the model against the large action space and non-stationarity problems. We also employ the distillation-based inference optimization.
- We evaluate the framework with several embodied environments and embedded devices, demonstrating its robustness and adaptability in terms of time-sensitive inference performance upon a wide range of tasks and operational conditions.

2. Related Work

Embodied AI. Achieving an embodied agent requires learning complex and diverse tasks, as well as adapting to a constantly changing, real-world environment. Many researchers focused on complex tasks in embodied environments, including object navigation [3, 5, 8, 33, 37], and embodied question and answering [6, 24, 31, 41], as well as model transfer from simulation to deployment environments [11, 17, 19]. Specifically, Gordon et al. [11] introduced SplitNet, decoupling visual perception and policy learning. By decomposing the network architecture into a visual encoder and a task decoder, it allows for rapid adaptation to new domains and vision tasks. Li et al. [19] also presented a learning framework that can adjust a learned policy to the target environment that differs from the training environment, utilizing unlabeled data from the target. While sharing the similar goal to adapt to different embodied environment conditions with the prior works, we focus on dynamic model adaptation to both time constraints and device limitations in the context of multi-task policy learning and inference.

Real-time model inference. Several works have been introduced in the realm of real-time model inference [1, 2, 14, 18, 23, 32, 38]. Specifically, Cai et al. [2] explored the trade-off between model performance and inference efficiency, by selecting certain nodes of the network and some of various filter-size CNN layers. Li et al. [18] introduced DS-Net, where weight ratios within each convolution neural network determine the network slicing for optimized inference. Unlike these works that enable the model adaptation for a given static condition, our work considers instance-wise operational conditions that can be given as input to the model for adaptation.

Model adaptation. Various works for quickly adapting a model to different environment features have been presented; e.g., for task difficulty levels [4, 15, 18, 30, 35], unseen tasks [9, 10, 20, 22, 40], and embodied properties [25, 27]. Wang et al. [35] proposed a framework that selectively skips CNN layers and channels within layers

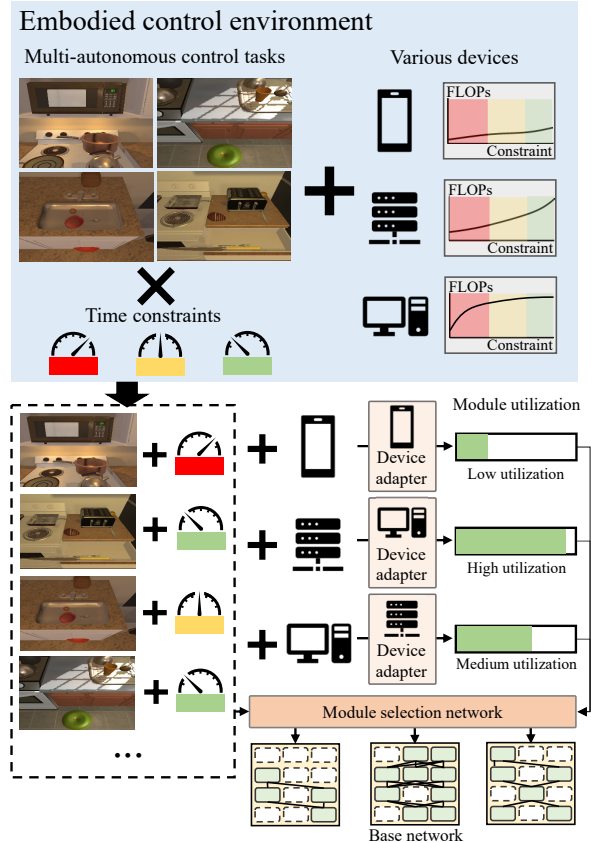


Figure 1. Overall Architecture

based on the task difficulty. They also exploited an early exit mechanism for resource-constrained inference. Han et al. [13] introduced a training algorithm and architecture that can adjust the latency by pixel-wise masking of CNNs, employing the latency prediction model and dynamic adjustment of hyperparameters for different devices. In line with these dynamic model adaptation works, we employ a modular network architecture and learning algorithm specifically designed for embodied control systems with multiple tasks, time constraints, and different device specifications.

3. Approach

3.1. Problem Formulation

We consider reinforcement learning (RL) for multi-autonomous control tasks in embodied environments. As an individual task is formulated as a single Markov decision process (MDP) \mathcal{M} , a multi-task MDP is equal to a family of MDPs $\{\mathcal{M}_i = (\mathcal{S}, \mathcal{A}, \mathcal{P}_i, R_i, \gamma)\}_i$. \mathcal{S} is a set of states, \mathcal{A} is a set of actions, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is a transition probability, $R : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is a reward function, and γ is a discount factor. In multi-task RL, task information is used to reformulate a family of MDPs in a single MDP. Accordingly, given a task index $i \in I$, such a reformulated

MDP can be represented as $(S \times I, \mathcal{A}, \mathcal{P}_I, R_I, \gamma)$ where $\mathcal{P}_I((s, i), a) = \mathcal{P}_i(s, a)$ and $R_I((s, i), a) = R_i(s, a)$. Furthermore, we consider time constraints in multi-task RL. Thus, a set of MDPs is represented as

$$\{\mathcal{M}_{i,c} = (S, \mathcal{A}, \mathcal{P}_{i,c}, R_{i,c}, \gamma)\}_{i,c} \quad (1)$$

where $c \in C$ is a time constraint. In this multi-task RL with time constraints, an embodied agent is learned to not only satisfy the time constraints but also maximize the cumulative discounted rewards for devices D where it can be deployed. Accordingly, the learning objective is to find the optimal policy π^* such as

$$\operatorname{argmax}_{\pi} \mathbb{E}_{I \times C} \left[\sum_{t=0}^N \gamma^t R_{i,c}(s_t, \pi(s_t)) \right], T_D(\pi) \leq c \quad (2)$$

where $T_D(\pi)$ represents the inference time on a device D .

3.2. Overall Approach

As illustrated in Figure 1, we address the problem of multi-task RL with time constraints for embodied control, by employing the dynamic module selection in a multi-task modular network. Our framework MoDeC includes three components: a modular base network for learning diverse tasks, a module selection network for performing adaptive inference under given time constraints, and a device adapter for configuring the module utilization according to the resource availability of a specific target device.

To achieve an adaptive policy, our approach utilizes a modular base network structure. This structure is flexible, allowing for direct adjustment of the computational load (in FLOPs) through selective module activation. It empowers the system to effectively balance the trade-off between accuracy and inference time (delay), thus enabling the time-sensitive robust model inference. To implement this structure, we adopt the soft modularization technique [39], which dynamically determines the weight of the path between learning modules for given task information. We implement the module selection network to determine the modules of the base network for each inference. This facilitates instance-wise computational adaptability, by taking task information and module utilization as input. By joint learning with the base network in a multi-task environment, the module selection network learns to determine the effective combination of modules for a specific task under the accuracy and inference time trade-off. Finally, to adapt to time constraints for each device, the device adapter converts the constraints into tolerable module utilization. This allows MoDeC to directly use the constraints for inference. Each device adapter is tailored for its own target device through few-shot learning.

During the model deployment on a specific target device, the device adapter infers the appropriate module utilization that adheres to given constraints. Then, the module

utilization is used as input to the module selection network that determines the modules to use for each input instance (i.e., each visual state for a multi-task embodied RL agent). Individual instances contain different task information, and each can be combined with the module utilization so as to make effective decisions on the module selection.

4. Learning A Modular Network

We describe the joint learning procedure for the modular base network and the module selection network.

Modular base network. To achieve a multi-task RL model, we employ soft modularization [39], a composite structure with a modular network and a soft routing network. The modular network infers actions based on state s , and the soft routing network infers the weights of paths in the modular network based on both state s and task index τ . We add a module index set to use m as input to the base network. The base network uses only the modules specified in m at inference. For batch $\mathcal{B} = \{(s, a, \tau)_i\}_{i \leq n}$ from replay buffer $\mathcal{D}_{\text{base}}$, we obtain a pre-trained base network $\bar{\pi}_{\text{base}}$ by optimizing multi-task loss $\mathcal{L}_{\text{MTRL}}$ defined as

$$\mathbb{E}_{\mathcal{B}} [w_{\tau} * (\alpha_{\tau} \log \bar{\pi}_{\text{base}}(a|s, \tau, m_{\text{full}}) - Q(s, a))]. \quad (3)$$

α_{τ} is a temperature parameter of entropy for each task τ , $m_{\text{full}} = [1, 1, 1, \dots, 1]$ represents selected modules, and Q is the learned Q-function. To promote multi-task RL in the set of tasks, we adjust the learning speed for each task by adjusting the scale of loss using weights calculated as

$$w_{\tau} = \frac{\exp(-\alpha_{\tau})}{\sum_{\tau \in \mathcal{T}} \exp(-\alpha_{\tau})}. \quad (4)$$

Joint learning. To enable an embodied agent to quickly adapt to time constraints, we jointly optimize the pre-trained base network and the module selection network. The module selection network infers module selection m based on the state s , task-specific information τ , and the number of modules to use K . The combination of the base network and module selection network can adjust the inference time by taking the number of modules to use as input and partially activating the modules of the base network.

In the process of joint learning, two problems arise: (1) the combinatorial optimization problem specific to module selection, and (2) the non-stationarity in concurrent RL training. The former comes from an exponentially large action space in module selection, significantly degrading the performance and learning efficiency. For instance, with 16 modules, potential combinations reach approximately 10^5 , complicating RL exploration and increasing sample amounts [26]. Furthermore, in joint learning, the interaction between the base network and the module selection network leads to a non-stationary learning environment [29].

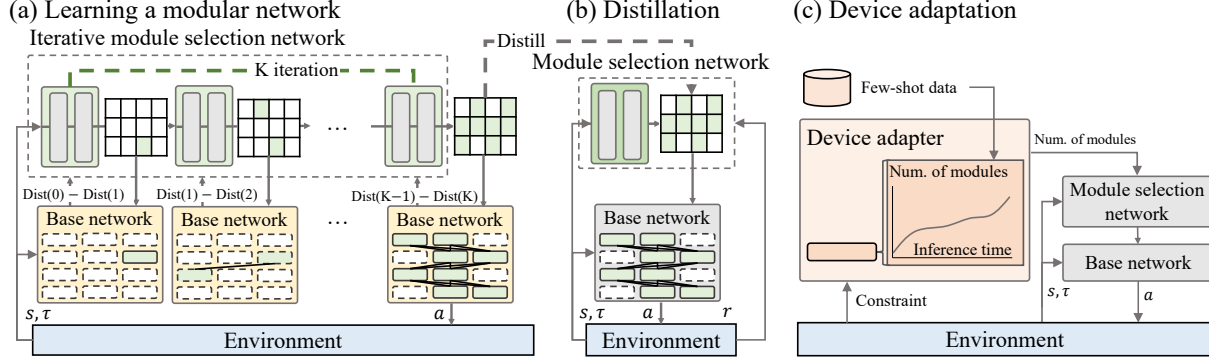


Figure 2. Learning Procedure of MoDeC. On the left side of the figure, the base network and the iterative module selection network are jointly leaned through a reward function R_{ims} . The iterative module selection network then distills into a single-step decision module selection network, as shown in the middle side. Finally, as depicted on the right side, the device adapter utilizes few-shot samples to associate the inference time with the number of modules (module utilization), effectively transforming the constraint representation into a specific number of modules to use for different devices.

Algorithm 1 Joint learning procedure

Multi-task environment env , Replay buffer $\mathcal{D}_{\text{base}}, \mathcal{D}_{\text{ims}} = \emptyset$
 Number of modules in base network N , timesteps t
 Learning rate $\lambda_{\text{base}}, \lambda_{\text{ims}}$, Pre-trained base network $\bar{\pi}_{\text{base}}$
 Base network π_{base} , Iterative module selection network π_{ims}

loop

$t = 0, s_0, \tau_0 = env.reset(), \pi_{\text{base}} \leftarrow \bar{\pi}_{\text{base}}$

$K \sim \text{Uniform}(\{1, \dots, N\})$

loop

$m_{0:0} = [0, \dots, 0]$

for $i = 1, \dots, K$ **do**

$\hat{m}_i = \pi_{\text{ims}}(s_t, \tau_t, K, m_{0:i-1})$

$m_{0:i} = m_{0:i-1} + \hat{m}_i$

$r_i = R_{\text{ims}}(\pi_{\text{base}}, s_t, \tau_t, m_{0:i})$ using (6)

$\mathcal{D}_{\text{ims}} \leftarrow \mathcal{D}_{\text{ims}} \cup \{(s_t, \tau_t, K, m_{0:i-1}), \hat{m}_i, r_i\}$

end for

$\pi_{\text{ims}} \leftarrow \pi_{\text{ims}} - \lambda_{\text{ims}} \cdot \nabla \mathcal{L}_{\text{IMS}}$ using (9)

$a_t = \pi_{\text{base}}(s_t, \tau_t, m_{0:K})$

$s_{t+1}, \tau_{t+1}, r_t = env.step(a_t)$

$\mathcal{D}_{\text{base}} \leftarrow \mathcal{D}_{\text{base}} \cup \{(s_t, a_t, r_t, s_{t+1})\}$

$\pi_{\text{base}} \leftarrow \pi_{\text{base}} - \lambda_{\text{base}} \cdot \nabla (\mathcal{L}_{\text{MTRL}} + \mathcal{L}_{\text{RG}})$ using (3), (8)

$t = t + 1$

end loop

end loop

Module selection network. To address the combinatorial optimization problem, we incorporate an iterative decision making procedure into module selection. The iterative module selection network, denoted as π_{ims} , operates by sequentially choosing an individual module across a total of K iterations. This achieves the selection of K modules represented as a binary vector $m \in \{0, 1\}^N$, where 1 indicates a selected module and 0 denotes a non-selected one. π_{ims} maps s, τ , and the cumulative sum of module selections $m_{0:i-1}$ until the $(i-1)^{\text{th}}$ iteration to the i^{th} individual mod-

ule selection denoted as $\hat{m}_i \in \{0, 1\}^N$. Thus, the module selection m for K modules is inferred by

$$m = \sum_{i=1}^K \pi_{\text{ims}}(s, \tau, K, m_{0:i-1}). \quad (5)$$

To train π_{ims} , we directly evaluate each selection, leveraging a reward function based on the similarity in actions inferred by the base network π_{base} .

Given an action inferred through utilizing the entire modules, the reward function R_{ims} is defined from the difference in distance between $\pi_{\text{base}}(s, \tau, m_{\text{full}})$ and $\pi_{\text{base}}(s, \tau, m_{0:i})$ subsequent to the previous module selection step:

$$R_{\text{ims}}(\pi_{\text{base}}, s, \tau, m_{0:i}) = \text{Dist}(i-1) - \text{Dist}(i) \quad (6)$$

where $\text{Dist}(i) = \|\pi_{\text{base}}(s, \tau, m_{\text{full}}) - \pi_{\text{base}}(s, \tau, m_{0:i})\|$. This reward function not only accelerates the learning of the iterative module selection network but also minimizes the regret bounds of the actions generated by the base network. When representing R for given task τ as an L -Lipschitz function, we can obtain the upper bound of the difference of rewards in a multi-task environment.

$$|R(s, \pi_{\text{base}}(s, \tau, m_{\text{full}})) - R(s, \pi_{\text{base}}(s, \tau, m_{0:i}))| \leq L \cdot \text{Dist}(i) \quad (7)$$

Thus, by minimizing $\text{Dist}(K)$, the difference in rewards in Eq. (6) is also minimized. This allows the actions inferred using a subset of modules to closely approximate the optimal reward.

To mitigate performance drops caused by the non-stationary problems, we avoid dramatic changes in actions between the pre-trained base network and the fine-tuned base network by using a regularization loss. Let $\mathcal{B} =$

$\{(s, \tau, K, m_{0:t_i-1}), \hat{m}_{t_i}, r_i\}_{i \leq n}$ be a sample batch from replay buffer \mathcal{D}_{ims} . The regularization loss \mathcal{L}_{RG} is defined as

$$\mathbb{E}_{\mathcal{B}} [\|\bar{\pi}_{\text{base}}(s, \tau, m_{\text{full}}) - \pi_{\text{base}}(s, \tau, m_{0:t_i})\|] \quad (8)$$

where $\bar{\pi}_{\text{base}}$ is the pre-trained base network without additional learning. The module selection m_{full} implies that whole modules of the base network are selected. Furthermore, we combine Reptile [28], a meta-RL algorithm, with REINFORCE [36]. The policy loss for the iterative module selection network \mathcal{L}_{IMS} is defined as

$$\mathbb{E}_{\mathcal{B}} \left[\sum_{i=1}^n \log \pi_{\text{ims}}(\hat{m}_{t_i} | s_i, \tau, K, m_{0:t_i-1}) \cdot G_i \right]. \quad (9)$$

Here, G_i is the discounted cumulative sum of rewards at timesteps i . Given that π_{ims}^k is the result of k updates from π_{ims} with the sample batch relative to the current π_{base} , the update for π_{ims} is executed with the step-size parameter ϵ .

$$\pi_{\text{ims}} \leftarrow \pi_{\text{ims}} + \epsilon(\pi_{\text{ims}}^k - \pi_{\text{ims}}) \quad (10)$$

The joint learning procedure of the base network and the iterative module selection network is illustrated on the left side of Figure 2, with details provided in Algorithm 1.

5. Distillation-based Optimization

We describe two schemes tailored for device-specific adaptation, the knowledge distillation for the module selection network and the few-shot learning for the device adapter. To enhance the efficiency of the module selection network, we reconstruct it with single-step inference through knowledge distillation. While the iterative module selection network shows superior performance in the large action space, its inference often incurs excessive delays and computational loads compared to the base network. The module selection network, denoted as π_{ms} , takes a $\mathcal{B} = \{(s, \tau, K)_i\}_{i < n}$ from replay buffer \mathcal{D}_{ms} as input in a single step. To train π_{ms} based on π_{ims} , we use $\mathcal{L}_{\text{MTRL}}$ in Eq. (3), with the knowledge distillation loss \mathcal{L}_{KD} defined as

$$\mathbb{E}_{\mathcal{B}} \left[\left\| \pi_{\text{ms}}(s, \tau, K) - \sum_{i=1}^K \pi_{\text{ims}}(s, \tau, K, m_{0:i-1}) \right\| \right]. \quad (11)$$

This distillation not only ensures the maintenance of the module selection performance but also considerably reduces the inference time with the module selection network.

6. Few-Shot Device Adaptation

To make our model adaptable under various time constraints for a specified device, we employ a device adapter that can manipulate the inference of the module selection network. This device adapter takes a time constraint, determining the

Algorithm 2 Distillation-based optimization

Multi-task environment env , device D , time-constraint c
Number of modules in base network N , timesteps t
Base network π_{base} , Iterative module selection network π_{ims}
Module selection network π_{ms} , Device adapter π_{da}
Inference time T_D for device D
Dataset $\mathcal{D}_{\text{ms}}, \mathcal{D}_{\text{da}} = \emptyset$, Learning rate $\lambda_{\text{ms}}, \lambda_{\text{da}}$
/ Distillation-based optimization of π_{ms} */*
loop
 $t = 0, s_0, \tau_0 = env.reset()$
 $K \sim \text{Uniform}(\{1, \dots, N\})$
loop
 $m_t = \pi_{\text{ms}}(s_t, \tau_t, K), a_t = \pi_{\text{base}}(s_t, \tau_t, m_t)$
 $m_{0:K} = \sum_{i=0}^K \pi_{\text{ims}}(s_t, \tau_t, K, m_{0:i})$
 $s_{t+1}, \tau_{t+1}, r_t = env.step(a)$
 $r_t = r_t - \|m_t - m_{0:K}\|_2$
 $\mathcal{D}_{\text{ms}} \leftarrow \mathcal{D}_{\text{ms}} \cup \{(s_t, \tau_t, K), m_t, r_t, (s_{t+1}, \tau_{t+1}, K)\}$
 $\pi_{\text{ms}} \leftarrow \pi_{\text{ms}} - \lambda_{\text{ms}} \cdot \nabla(\mathcal{L}_{\text{MTRL}} + \mathcal{L}_{\text{KD}})$ using (3), (11)
 $t = t + 1$
end loop
end loop
/ Few-shot adaptation through device adapter π_{da} */*
loop
 $t = 0, s_0, \tau_0 = env.reset()$
 $K \sim \text{Uniform}(\{1, \dots, N\})$
loop
 $c_t, a_t = T_D(\pi_{\text{base}}(s_t, \tau_t, \pi_{\text{ms}}(s_t, \tau_t, K)))$
 $s_{t+1}, \tau_{t+1}, r_t = env.step(a_t)$
 $\mathcal{D}_{\text{da}} \leftarrow \mathcal{D}_{\text{da}} \cup \{(K, c_t)\}$
 $\pi_{\text{da}} \leftarrow \pi_{\text{da}} - \lambda_{\text{da}} \cdot \nabla \mathcal{L}_{\text{DA}}$ using (12)
 $t = t + 1$
end loop
end loop

number of modules to use, ensuring that it does not violate the constraint. By using the adapter, time constraints are directly grounded as values within the network, enabling MoDeC to perform the constraint-aware inference.

To train the device adapter π_{da} , we use a pre-trained base network π_{base} and distilled module selection model π_{ms} with a loss function specifically designed to accommodate the constraints of the current device. For a given device, we collect a model inference dataset \mathcal{D}_{da} and sample batch denoted as $\mathcal{B} = \{(K, c)_i\}_{i < n}$, where K is the number of modules to use and c is the inference time of MoDeC when using only K modules. The device adapter is optimized by \mathcal{L}_{DA} .

$$\mathbb{E}_{\mathcal{B}} [\|\pi_{\text{da}}(c) - K\| \cdot (1 - p)] \quad (12)$$

Here, p represents the penalty weights applied when the device adapter predicts the number of modules exceeding K ; otherwise, p is set to 0.

The procedure of distillation-based optimization and few-shot device adaptation is illustrated on the right side of Figure 2, with details provided in Algorithm 2.

Table 1. Performance for Meta-world single task in the success rate with 95% confidence intervals: the best performance is in bold.

Device	Constraint	DRNet		D2NN		DS-Net		RL-AA		MoDeC	
		Success rate	FLOPs	Success rate	FLOPs	Success rate	FLOPs	Success rate	FLOPs	Success rate	FLOPs
Orin	8 ms	-	-	26.7 ± 3.9%	62M	34.8 ± 11.7%	25M	35.4 ± 4.4%	149M	57.5 ± 9.9%	92M
	10 ms	18.0 ± 4.5%	190M	27.3 ± 3.5%	62M	30.5 ± 9.3%	159M	34.9 ± 5.0%	149M	65.0 ± 13.4%	151M
	12 ms	18.0 ± 4.5%	190M	32.0 ± 4.4%	110M	31.8 ± 12.7%	403M	41.2 ± 4.1%	355M	75.0 ± 9.4%	229M
	14 ms	38.1 ± 10.6%	364M	54.0 ± 6.9%	176M	31.5 ± 9.9%	567M	47.1 ± 5.0%	360M	74.5 ± 14.2%	254M
	16 ms	39.8 ± 11.7%	491M	68.0 ± 8.9%	284M	30.0 ± 8.1%	758M	47.3 ± 5.9%	365M	74.5 ± 14.2%	254M
Xavier	12 ms	-	-	27.3 ± 4.2%	62M	31.0 ± 9.5%	25.48M	33.4 ± 5.3%	149M	50.0 ± 7.5%	92M
	15 ms	18.0 ± 4.5%	190M	36.0 ± 6.8%	62M	30.0 ± 8.5%	78M	32.1 ± 4.9%	205M	56.0 ± 9.4%	151M
	18 ms	20.0 ± 6.7%	243M	38.2 ± 9.0%	62M	26.0 ± 6.9%	159M	41.3 ± 4.4%	355M	86.0 ± 9.7%	214M
	21 ms	20.0 ± 6.7%	243M	33.8 ± 5.9%	160M	29.3 ± 8.8%	403M	45.7 ± 8.5%	365M	80.0 ± 13.5%	254M
	24 ms	38.0 ± 10.7%	391M	69.6 ± 9.0%	284M	28.0 ± 7.4%	567M	48.0 ± 6.4%	365M	80.0 ± 13.5%	254M
Nano	40 ms	-	-	30.0 ± 7.5%	62M	32.0 ± 6.7%	78M	32.1 ± 6.1%	149M	40.0 ± 9.2%	114M
	46 ms	18.0 ± 4.5%	190M	24.1 ± 6.0%	82M	30.7 ± 4.0%	159M	40.4 ± 7.7%	355M	47.5 ± 13.2%	158M
	52 ms	37.9 ± 10.5%	320M	38.0 ± 12.5%	82M	28.7 ± 3.9%	403M	42.3 ± 5.3%	355M	72.5 ± 10.2%	214M
	58 ms	40.9 ± 10.3%	485M	38.7 ± 10.3%	180M	28.7 ± 4.5%	567M	46.1 ± 4.9%	365M	65.0 ± 9.2%	258M
	64 ms	42.0 ± 12.5%	539M	68.7 ± 5.5%	269M	26.7 ± 3.9%	758M	46.9 ± 4.4%	365M	65.0 ± 9.2%	258M

7. Evaluation

7.1. Environments and Devices




Figure 3. Environments

Meta-world. We use the MT10 benchmark (i.e., 10 different control tasks) in Meta-world [42], where each task is given a specific manipulation objective such as opening a door or closing a window. We compare the performance of robot manipulation tasks under time constraints.

CARLA. To demonstrate mission-critical scenarios where the inference time is of critical importance, we use the autonomous driving simulator CARLA [7]. Models are trained for autonomous driving tasks with vision-based states at a multi-task configuration with 12 different maps.

AI2THOR. We use AI2THOR [16], where an agent navigates the map with egocentric vision states, placing various objects to complete a rearrangement task. The simulation environments are represented in Figure 3. In our evaluation, we test several embedded devices, each with distinct resources and computational capabilities. The devices include Nvidia Jetson Nano (Nano), Nvidia Jetson Xavier NX 8GB (Xavier), and Nvidia Jetson AGX Orin 32GB (Orin), with the Nano being the least powerful, followed by the Xavier and the Orin being the most powerful. By testing on these devices with varying levels of capabilities, we can better understand how our framework adapts to different resource limitations. This is crucial in embodied AI, where deployment environments can greatly vary in terms of available

computational resources. The detailed device specifications are in Table 2.

Table 2. Device Specification

Device	Performance	CPU Max Freq.	GPU Max Freq.	Memory
Orin	275 TOPs	2.2 GHz	1.3 GHz	32 GBs
Xavier	21 TOPs	1.9 GHz	1.1 GHz	8 GBs
Nano	472 GFLOPs	0.9 GHz	0.6 GHz	4 GBs

7.2. Comparisons

We use several dynamic model adaptation methods as baselines.

- Dynamic Routing Network (DRNet) [2] is a network comprising serially connected cells, each corresponding to a directed acyclic graph of nodes. It optimizes by learning to select paths between the nodes through a loss function that balances the inference time and performance. Unlike MoDeC which adapts a single model to different conditions, we use individual networks specifically learned for each constraint condition. We consider DRNet as a baseline for dynamic (adaptive) inference models.
- Dynamic Deep Neural Networks (D2NN) [23] is a modular neural network that exploits the accuracy-efficiency trade-off. It exploits RL in module selection, using the rewards calculated according to performance and inference time. Unlike MoDeC, we use individual networks specifically learned for each constraint condition. We consider D2NN as an RL baseline tailored for trade-off conditions.
- Dynamic Slimmable Network (DS-Net) [18] is a dynamic network model, in which an internal gater network determines the weight ratio for convolution neural layers to slice the network for inference. To align with our problem formulation, we modify the gater network (taking the ratio as input) in a way of conducting instance-wise dynamic inference, similar to our approach. In our comparison, DS-Net serves as a baseline for adaptive models that handle multiple constraints within a single policy.
- RL via Asymmetric Architecture (RL-AA) [4] is a hi-

Table 3. Performance for Meta-world multi-task in the success rate with 95% confidence intervals

Device	Constraint	DRNet		D2NN		DS-Net		RL-AA		MoDeC	
		Success rate	FLOPs	Success rate	FLOPs	Success rate	FLOPs	Success rate	FLOPs	Success rate	FLOPs
Orin	8 ms	-	-	20.0 ± 6.7%	92M	11.7 ± 7.1%	25M	15.7 ± 6.9%	149M	27.1 ± 8.3%	92M
	10 ms	12.0 ± 4.2%	119M	18.0 ± 5.6%	151M	11.6 ± 4.4%	159M	15.3 ± 7.1%	149M	30.3 ± 10.1%	151M
	12 ms	26.7 ± 5.4%	488M	17.0 ± 4.8%	229M	23.3 ± 8.5%	403M	18.5 ± 8.6%	355M	34.9 ± 9.7%	229M
	14 ms	26.7 ± 6.1%	488M	51.9 ± 10.4%	254M	16.7 ± 8.7%	567M	26.1 ± 9.1%	355M	53.0 ± 6.8%	254M
	16 ms	30.1 ± 4.3%	636M	54.1 ± 8.4%	254M	25.0 ± 5.7%	758M	28.3 ± 7.1%	365M	53.0 ± 6.8%	254M
Xavier	12 ms	-	-	20.4 ± 6.7%	92M	18.3 ± 7.9%	25M	14.0 ± 7.8%	149M	26.7 ± 7.6%	92M
	15 ms	10.3 ± 3.2%	119M	23.5 ± 6.9%	151M	16.7 ± 5.4%	78M	14.9 ± 5.3%	149M	29.5 ± 10.1%	151M
	18 ms	19.8 ± 4.5%	488M	17.3 ± 4.8%	214M	18.3 ± 4.2%	159M	23.1 ± 9.2%	355M	28.1 ± 9.1%	214M
	21 ms	18.8 ± 3.7%	488M	56.8 ± 10.1%	254M	21.7 ± 10.3%	403M	28.1 ± 6.2%	365M	53.6 ± 5.9%	254M
	24 ms	21.2 ± 3.8%	636M	53.0 ± 9.0%	254M	15.0 ± 5.8%	567M	27.6 ± 6.5%	365M	53.6 ± 5.9%	254M
Nano	40 ms	-	-	23.5 ± 15.2%	114M	10.4 ± 8.9%	78M	16.1 ± 6.1%	149M	27.0 ± 7.5%	114M
	46 ms	10.0 ± 0.8%	119M	17.5 ± 8.0%	158M	9.9 ± 2.1%	159M	23.4 ± 7.0%	355M	23.0 ± 5.8%	162M
	52 ms	23.3 ± 5.4%	488M	20.0 ± 18.3%	214M	13.3 ± 5.3%	267M	24.1 ± 6.5%	365M	28.1 ± 8.1%	217M
	58 ms	16.7 ± 4.9%	488M	52.3 ± 8.0%	258M	20.0 ± 9.4%	567M	27.4 ± 8.6%	365M	53.9 ± 6.0%	234M
	64 ms	23.3 ± 6.1%	636M	52.5 ± 17.9%	258M	20.0 ± 5.4%	758M	26.9 ± 6.9%	365M	62.1 ± 4.5%	254M

Table 4. Performance on CARLA in the success rate with 95% confidence intervals

Device	Constraint	DRNet		D2NN		DS-Net		RL-AA		MoDeC	
		Success rate	FLOPs	Success rate	FLOPs	Success rate	FLOPs	Success rate	FLOPs	Success rate	FLOPs
Orin	8 ms	0.0 ± 0.0%	16M	0.0 ± 0.0%	32M	0.0 ± 0.0%	3M	0.0 ± 0.0%	34M	8.3 ± 5.0%	31M
	10 ms	0.0 ± 0.0%	16M	0.0 ± 0.0%	36M	39.9 ± 9.0%	10M	0.0 ± 0.0%	34M	49.7 ± 9.1%	47M
	12 ms	70.1 ± 6.1%	27M	16.6 ± 6.8%	42M	71.1 ± 7.2%	36M	33.1 ± 5.9%	42M	83.1 ± 6.8%	60M
	14 ms	68.2 ± 5.4%	48M	41.6 ± 8.9%	57M	66.7 ± 8.6%	76M	75.6 ± 4.9%	72M	79.2 ± 6.7%	73M
	16 ms	70.3 ± 5.6%	48M	83.3 ± 6.7%	72M	76.7 ± 7.7%	132M	77.1 ± 5.1%	72M	85.8 ± 5.8%	73M
Xavier	12 ms	0.0 ± 0.0%	16M	0.0 ± 0.0%	36M	0.0 ± 0.0%	3M	0.0 ± 0.0%	34M	12.3 ± 6.9%	35M
	15 ms	0.0 ± 0.0%	16M	0.0 ± 0.0%	36M	35.5 ± 5.0%	10M	0.0 ± 0.0%	34M	41.7 ± 7.7%	47M
	18 ms	52.1 ± 13.1%	27M	27.8 ± 8.6%	48M	77.1 ± 9.1%	36M	36.1 ± 5.1%	42M	80.8 ± 6.7%	59M
	21 ms	71.1 ± 7.3%	48M	41.7 ± 9.4%	54M	75.2 ± 7.9%	76M	58.1 ± 4.8%	58M	79.1 ± 7.3%	71M
	24 ms	72.2 ± 8.6%	48M	89.3 ± 4.4%	72M	83.3 ± 6.8%	132M	76.4 ± 5.9%	72M	83.3 ± 6.7%	73M
Nano	40 ms	0.0 ± 0.0%	16M	0.0 ± 0.0%	16M	0.0 ± 0.0%	3M	0.0 ± 0.0%	34M	10.4 ± 4.1%	35M
	46 ms	0.0 ± 0.0%	16M	0.0 ± 0.0%	36M	0.0 ± 0.0%	3M	32.7 ± 5.4%	42M	35.5 ± 10.3%	43M
	52 ms	70.7 ± 5.6%	27M	51.2 ± 6.6%	54M	23.1 ± 15.3%	10M	34.5 ± 4.6%	42M	77.1 ± 5.3%	56M
	58 ms	69.1 ± 7.1%	48M	83.5 ± 7.0%	72M	59.1 ± 11.2%	32M	74.1 ± 4.3%	72M	82.0 ± 5.7%	71M
	64 ms	72.9 ± 5.3%	48M	85.8 ± 6.1%	72M	65.5 ± 6.7%	36M	75.3 ± 5.6%	72M	81.9 ± 6.1%	73M

Table 5. Performance on AI2THOR in the success rate

Device	Constraint	DS-Net		MoDeC	
		Seen	Unseen	Seen	Unseen
Orin	8 ms	98.7%	5.8%	97.9%	21.2%
	10 ms	96.3%	5.2%	98.0%	45.4%
	12 ms	98.5%	4.8%	97.5%	41.7%
	14 ms	96.3%	6.0%	98.2%	71.3%
	16 ms	97.5%	7.4%	97.9%	81.2%
Xavier	12 ms	97.1%	6.2%	97.2%	20.7%
	15 ms	95.7%	8.0%	98.7%	44.0%
	18 ms	98.1%	7.4%	97.5%	45.2%
	21 ms	97.7%	7.0%	98.0%	69.3%
	24 ms	98.0%	8.2%	98.1%	79.1%

erarchical policy to dynamically adjust the module usage. The low-level policy consists of two models, each with a small and large scale, and the high-level policy determines which policy to use. To adapt to various time-constrained conditions, we include a wider range of low-level policies, each with varying inference time. We use RL-AA as a baseline for resource-adaptive RL methods.

7.3. Adaptation Performance

Meta-world Single-task. For 5 individual tasks in MT10, Table 1 shows the performance under various time con-

straints (in the column of “Constraint”), achieved by our MoDeC and the baselines (DRNet, D2NN, DS-Net, RL-AA). Specifically, we evaluate the average success ratio and computation load (in FLOPs) within the constraint violation rate of 1% for 3 different devices (in the column of “Constraint”). As shown, MoDeC achieves superior performance for most configurations. Compared to D2NN, the most competitive baseline, MoDeC achieves a 25.1% gain. While sharing a common base network structure with D2NN, MoDeC shows better performance, as it employs the iterative module selection and distillation. More importantly, D2NN needs to be retrained for each configuration (i.e., each constraint and device setting). MoDeC achieves this performance superiority across different configurations, using only a single model without retraining, demonstrating its adaptation capabilities to different time and resource constraints.

Meta-world Multi-task. Table 3 compares the performance of the MT10 multi-task. MoDeC demonstrates consistently its performance superiority, achieving an average performance gain of 5.7% over D2NN, the most competitive baseline. This specifies the adaptation capabilities of MoDeC, achieved not only through module selection but

also through multi-task learning for the base network.

CARLA. Table 4 shows the performance for autonomous driving tasks across 12 different maps in CARLA, where delayed inference often degrades the performance and poses risks; we implement such a strategy that upon a constraint violation (i.e., inference delay), the action at the previous timestep is reused. MoDeC shows 14.4% higher performance than DS-Net, which is the most competitive comparison in this experiment. Due to the direct impacts of constraint violations in CARLA, the adaptive inference is more beneficial, compared to the Meta-World tasks. This leads to better performance by the methods capable of constraint-aware inference, such as ours and DS-Net. DS-Net shows a significant performance drop in Nano, which is a small memory device; DS-Net requires a large computation load per single layer, unlike ours.

AI2THOR. Table 5 compares the performance for AI2THOR’s complex navigation tasks, where “Seen” refers to initial object positions encountered during training and “Unseen” refers to those not during training. Both MoDeC and DS-Net perform well in the seen configurations, but MoDeC demonstrates significantly better performance in the unseen configurations, showing a performance gap ranging from 15.4% to 70.8%. In MoDeC, the soft modularization combined with module selection facilitates effective module combinations for different tasks and constraints, rendering robust performance in unseen configurations.

7.4. Ablation Study

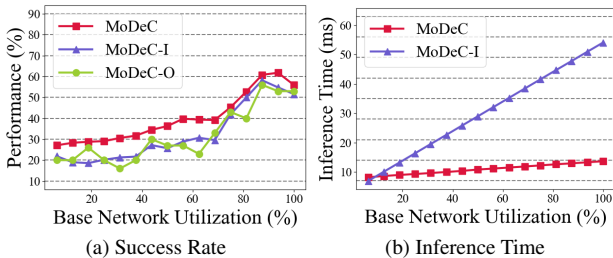


Figure 4. Effect of Distillation

Distillation. Figure 4 clarifies the effects of distillation, where MoDeC-I denotes a MoDeC variant without distillation, which adopts only the iterative module selection network; MoDeC-O denotes another variant, which directly learns the single-step model selection (without distillation). As shown, there is a significant difference in inference time between MoDeC and MoDeC-I as the module utilization increases in (b), while MoDeC (with distillation) achieves higher performance in (a). This is because the iterative module selection network is learned to infer as closely as possible to the original action under a limited module utilization, excluding environment rewards. When using both environment rewards and action distance, we observe a decline in

the performance of the iterative module selective network. Due to changes in the environment reward, the actions of the iterative module selection network are not properly evaluated. The performance decline in MoDeC-O stems from learning the module selection, which requires extensive exploration, yet is difficult in a single step without distillation.




Figure 5. Effect of Base Network Architecture

Base network architecture. Figure 5 illustrates the effects of the modular base network architecture, where its number of layers can be configured differently. In our implementation, the base network has 4 layers with 4 modules per layer, each represented as 4×4 in the figure. We compare this with other variants, 2×8 and 8×2 , in Meta-World. MoDeC shows the best performance by 4×4 , which is a hyperparameter in the base network architecture.

8. Conclusions

In this work, we presented MoDeC, which allows embodied agents to effectively adapt to time constraints on different target devices. The modular multi-task learning in MoDeC enables adaptive inference to a wide range of operational conditions including device resources, time constraints, and task specifications of embodied agents’ multi-task settings, by dynamically adjusting the inference within a single model to satisfy the operational requirement. Through experiments with manipulation, autonomous driving, and object navigation scenarios of embodied agents, we verified that MoDeC is capable of handling those control tasks through rapid model adaptation to various operational conditions that can change over time. Our future work is to tackle the challenge of learning complex constraints from instructions, including safety and resource limitations.

Acknowledgements

We would like to thank anonymous reviewers for their valuable comments and suggestions. This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2022-0-01045, 2022-0-00043, 2021-0-00875, 2020-0-01821, 2019-0-00421) and by the National Research Foundation of Korea (NRF) grant funded by the MSIT (No. RS-2023-00213118).

References

- [1] Tolga Bolukbasi, Joseph Wang, Ofer Dekel, and Venkatesh Saligrama. Adaptive neural networks for efficient inference. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pages 527–536. PMLR, 2017. 1, 2
- [2] Shaofeng Cai, Yao Shu, and Wei Wang. Dynamic routing networks. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3588–3597, 2021. 1, 2, 6
- [3] Tommaso Campari, Paolo Eccher, Luciano Serafini, and Lamberto Ballan. Exploiting scene-specific features for object goal navigation. In *Proceedings of the 16th European Conference on Computer Vision (ECCV)*, pages 406–421. Springer, 2020. 2
- [4] Chin-Jui Chang, Yu-Wei Chu, Chao-Hsien Ting, Hao-Kang Liu, Zhang-Wei Hong, and Chun-Yi Lee. Reducing the deployment-time inference control costs of deep reinforcement learning agents via an asymmetric architecture. In *Proceedings of the 38th IEEE International Conference on Robotics and Automation (ICRA)*, pages 4762–4768. IEEE, 2021. 2, 6
- [5] Devendra Singh Chaplot, Dhiraj Prakashchand Gandhi, Abhinav Gupta, and Russ R Salakhutdinov. Object goal navigation using goal-oriented semantic exploration. In *Proceedings of the 34th Conference on Neural Information Processing Systems (NeurIPS)*, pages 4247–4258, 2020. 2
- [6] Abhishek Das, Samyak Datta, Georgia Gkioxari, Stefan Lee, Devi Parikh, and Dhruv Batra. Embodied question answering. In *Proceedings of the 29th IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–10, 2018. 2
- [7] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. In *Proceedings of the 1st Conference on Robot Learning (CoRL)*, pages 1–16. PMLR, 2017. 1, 6
- [8] Heming Du, Xin Yu, and Liang Zheng. Learning object relation graph and tentative policy for visual navigation. In *Proceedings of the 16th European Conference on Computer Vision (ECCV)*, pages 19–34. Springer, 2020. 2
- [9] Daniel Fried, Ronghang Hu, Volkan Cirik, Anna Rohrbach, Jacob Andreas, Louis-Philippe Morency, Taylor Berg-Kirkpatrick, Kate Saenko, Dan Klein, and Trevor Darrell. Speaker-follower models for vision-and-language navigation. In *Proceedings of the 32nd Conference on Neural Information Processing Systems (NeurIPS)*, 2018. 2
- [10] Xiaofeng Gao, Qiaozi Gao, Ran Gong, Kaixiang Lin, Govind Thattai, and Gaurav S Sukhatme. Dialfred: Dialogue-enabled agents for embodied instruction following. *IEEE Robotics and Automation Letters*, 7(4):10049–10056, 2022. 2
- [11] Daniel Gordon, Abhishek Kadian, Devi Parikh, Judy Hoffman, and Dhruv Batra. Splitnet: Sim2sim and task2task transfer for embodied visual navigation. In *Proceedings of the 18th IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1022–1031, 2019. 2
- [12] Siddhant Haldar and Lerrel Pinto. Polytask: Learning unified policies through behavior distillation. *arXiv preprint arXiv:2310.08573*, 2023. 1
- [13] Yizeng Han, Zhihang Yuan, Yifan Pu, Chenhao Xue, Shiji Song, Guangyu Sun, and Gao Huang. Latency-aware spatial-wise dynamic networks. 35:36845–36857, 2022. 2
- [14] Weizhe Hua, Yuan Zhou, Christopher M De Sa, Zhiru Zhang, and G Edward Suh. Channel gating neural networks. In *Proceedings of the 33rd Conference on Neural Information Processing Systems (NeurIPS)*, 2019. 2
- [15] Gao Huang, Danlu Chen, Tianhong Li, Felix Wu, Laurens Van Der Maaten, and Kilian Q Weinberger. Multi-scale dense networks for resource efficient image classification. *arXiv preprint arXiv:1703.09844*, 2017. 2
- [16] Eric Kolve, Roozbeh Mottaghi, Winson Han, Eli VanderBilt, Luca Weihs, Alvaro Herrasti, Matt Deitke, Kiana Ehsani, Daniel Gordon, Yuke Zhu, et al. Ai2-thor: An interactive 3d environment for visual ai. *arXiv preprint arXiv:1712.05474*, 2017. 1, 6
- [17] Klemen Kotar and Roozbeh Mottaghi. Interactron: Embodied adaptive object detection. In *Proceedings of the 33rd IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14860–14869, 2022. 2
- [18] Changlin Li, Guangrun Wang, Bing Wang, Xiaodan Liang, Zhihui Li, and Xiaojun Chang. Dynamic slimmable network. In *Proceedings of the 23rd IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8607–8617, 2021. 1, 2, 6
- [19] Juncheng Li, Xin Wang, Siliang Tang, Haizhou Shi, Fei Wu, Yueting Zhuang, and William Yang Wang. Unsupervised reinforcement learning of transferable meta-skills for embodied navigation. In *Proceedings of the 31st IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12123–12132, 2020. 2
- [20] Jialu Li, Hao Tan, and Mohit Bansal. Envedit: Environment editing for vision-and-language navigation. In *Proceedings of the 33rd IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 15407–15417, 2022. 2
- [21] Chen Liang, Simiao Zuo, Qingru Zhang, Pengcheng He, Weizhu Chen, and Tuo Zhao. Less is more: Task-aware layer-wise distillation for language model compression. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, pages 20852–20867. PMLR, 2023. 1
- [22] Chong Liu, Fengda Zhu, Xiaojun Chang, Xiaodan Liang, Zongyuan Ge, and Yi-Dong Shen. Vision-language navigation with random environmental mixup. In *Proceedings of the 19th IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1644–1654, 2021. 2
- [23] Lanlan Liu and Jia Deng. Dynamic deep neural networks: Optimizing accuracy-efficiency trade-offs by selective execution. In *Proceedings of the 34nd AAAI Conference on Artificial Intelligence*, 2018. 1, 2, 6
- [24] Haonan Luo, Guosheng Lin, Fumin Shen, Xingguo Huang, Yazhou Yao, and Hengtao Shen. Robust-eqa: robust learning for embodied question answering with noisy labels. *IEEE Transactions on Neural Networks and Learning Systems*, 2023. 2

- [25] Arjun Majumdar, Karmesh Yadav, Sergio Arnaud, Yecheng Jason Ma, Claire Chen, Sneha Silwal, Aryan Jain, Vincent-Pierre Berges, Pieter Abbeel, Jitendra Malik, et al. Where are we in the search for an artificial visual cortex for embodied intelligence? *arXiv preprint arXiv:2303.18240*, 2023. 2
- [26] Luke Metz, Julian Ibarz, Navdeep Jaitly, and James Davidson. Discrete sequential prediction of continuous actions for deep rl. *arXiv preprint arXiv:1705.05035*, 2017. 3
- [27] Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhinav Gupta. R3m: A universal visual representation for robot manipulation. *arXiv preprint arXiv:2203.12601*, 2022. 2
- [28] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018. 5
- [29] Shayegan Omidshafiei, Jason Pazis, Christopher Amato, Jonathan P How, and John Vian. Deep decentralized multi-task multi-agent reinforcement learning under partial observability. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pages 2681–2690. PMLR, 2017. 3
- [30] Dripta S Raychaudhuri, Yumin Suh, Samuel Schuster, Xiang Yu, Masoud Faraki, Amit K Roy-Chowdhury, and Manmohan Chandraker. Controllable dynamic multi-task architectures. In *Proceedings of the 33rd IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10955–10964, 2022. 2
- [31] Sinan Tan, Mengmeng Ge, Di Guo, Huaping Liu, and Fuchun Sun. Knowledge-based embodied question answering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 2
- [32] Surat Teerapittayanon, Bradley McDanel, and Hsiang-Tsung Kung. Branchynet: Fast inference via early exiting from deep neural networks. In *Proceedings of the 23rd International Conference on Pattern Recognition (ICPR)*, pages 2464–2469. IEEE, 2016. 2
- [33] Ayzaan Wahid, Austin Stone, Kevin Chen, Brian Ichter, and Alexander Toshev. Learning object-conditioned exploration using distributed soft actor critic. In *Proceedings of the 5th Conference on Robot Learning (CoRL)*, pages 1684–1695. PMLR, 2021. 2
- [34] Longguang Wang, Xiaoyu Dong, Yingqian Wang, Li Liu, Wei An, and Yulan Guo. Learnable lookup table for neural network quantization. In *Proceedings of the 33rd IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12423–12433, 2022. 1
- [35] Yue Wang, Jianghao Shen, Ting-Kuei Hu, Pengfei Xu, Tan Nguyen, Richard Baraniuk, Zhangyang Wang, and Yingyan Lin. Dual dynamic inference: Enabling more efficient, adaptive, and controllable deep inference. *IEEE Journal of Selected Topics in Signal Processing*, 14(4):623–633, 2020. 2
- [36] Ronald J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8:229–256, 1992. 5
- [37] Mitchell Wortsman, Kiana Ehsani, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. Learning to learn how to learn: Self-adaptive visual navigation using meta-learning. In *Proceedings of the 30th IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6750–6759, 2019. 2
- [38] Wenhan Xia, Hongxu Yin, Xiaoliang Dai, and Niraj K Jha. Fully dynamic inference with deep neural networks. *IEEE Transactions on Emerging Topics in Computing*, 10(2):962–972, 2021. 2
- [39] Ruihan Yang, Huazhe Xu, Yi Wu, and Xiaolong Wang. Multi-task reinforcement learning with soft modularization. In *Proceedings of the 34th conference on neural information processing systems (NeurIPS)*, pages 4767–4777, 2020. 1, 3
- [40] Felix Yu, Zhiwei Deng, Karthik Narasimhan, and Olga Russakovsky. Take the scenic route: Improving generalization in vision-and-language navigation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 920–921, 2020. 2
- [41] Licheng Yu, Xinlei Chen, Georgia Gkioxari, Mohit Bansal, Tamara L Berg, and Dhruv Batra. Multi-target embodied question answering. In *Proceedings of the 30th IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6309–6318, 2019. 2
- [42] Tianhe Yu, Deirdre Quillen, Zhanpeng He, Ryan Julian, Karol Hausman, Chelsea Finn, and Sergey Levine. Meta-world: A benchmark and evaluation for multi-task and meta reinforcement learning. In *Proceedings of the 4th conference on robot learning (CoRL)*, pages 1094–1100. PMLR, 2020. 1, 6