

Multi-Agent Behavior Retrieval: Retrieval-Augmented Policy Training for Cooperative Push Manipulation by Mobile Robots

So Kuroki^{*1}, Mai Nishimura¹, and Tadashi Kozuno¹

Abstract—Due to the complex interactions between agents, learning multi-agent control policy often requires a prohibitive amount of data. This paper aims to enable multi-agent systems to effectively utilize past memories to adapt to novel collaborative tasks in a data-efficient fashion. We propose the Multi-Agent Coordination Skill Database, a repository for storing a collection of coordinated behaviors associated with key vectors distinctive to them. Our Transformer-based skill encoder effectively captures spatio-temporal interactions that contribute to coordination and provides a unique skill representation for each coordinated behavior. By leveraging only a small number of demonstrations of the target task, the database enables us to train the policy using a dataset augmented with the retrieved demonstrations. Experimental evaluations demonstrate that our method achieves a significantly higher success rate in push manipulation tasks compared with baseline methods like few-shot imitation learning. Furthermore, we validate the effectiveness of our retrieve-and-learn framework in a real environment using a team of wheeled robots.

I. INTRODUCTION

When making decisions or engaging in collaborative tasks such as moving heavy loads [1] or playing team sports [2], we effectively tap into past experiences by retrieving relevant memories. By drawing on the wealth of lessons learned from past successes and failures, we gain valuable insights that enhance our coordination ability to work together more efficiently even in unknown situations. In this paper, we focus on endowing a team of mobile robots with such intelligent memory storage, which we refer to as a *Multi-Agent Coordination Skill Database*. The database, which serves as a repository of *coordination skills*, efficiently stores collaborative demonstrations in diverse settings. Collecting a large number of multi-robot demonstrations for every target configuration is extremely costly and impractical in real-world scenarios. In contrast, with only a few demonstrations of the target task as queries, our database enables us to retrieve and learn from the vast repository of coordination skills that are relevant to the target task.

In this study, we explore the potential of such a system for push manipulation tasks by a team of cooperative mobile robots. While cooperative push manipulation enables more dexterous and stable handling of a target object [3], [4], capturing complex interactions between agents and an object often requires either a prohibitive amount of training data [4] or a specific formulation for each setup [5], [6]. Given the growing interest in cooperative push manipulation

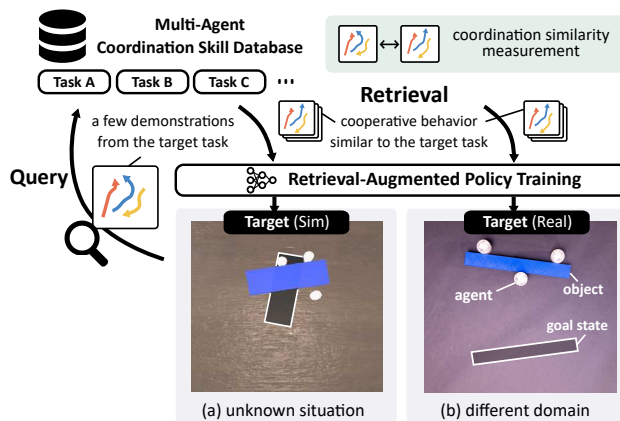


Fig. 1: Given a collection of unlabeled demonstrations from multiple tasks, our aim is to construct a coordination skill database that enables us to retrieve past experiences for different scenarios or domains (e.g., real-world robots).

across diverse fields such as factory logistics and delivery services, improving data efficiency is crucial for expanding its application and driving future advancements.

To achieve data-efficient learning for cooperative manipulation based on the skill database concept, this study focuses on the key research question: “How can past coordination behaviors of multiple agents be effectively retrieved and reused in new target tasks that require similar coordination?” While some pioneering works [7], [8] seek a similar direction of research that retrieves and learns from sub-optimal offline datasets, their studies are limited to single agents with image-based observations. Unlike single-agent scenarios, multi-agent push manipulations require sophisticated interactions among agents. This involves local cooperation with nearby agents to apply force to an object considering the object’s shape, while also making cooperative global decisions on the overall movement direction of the object by multiple robots, such as pushing or rotating. However, measuring the similarity between two cooperative demonstrations is a non-trivial problem. The local similarity observed between individual agents does not always conform to the global similarity among all agents. This discrepancy between local and global similarity poses a challenge for using these demonstrations in downstream policy training.

To address this question, we introduce a novel retrieve-and-learn framework based on the Multi-Agent Coordination Skill Database for push manipulations. Figure 1 depicts an overview of the proposed concept. We first train a Transformer-based coordination skill encoder that predicts

^{*}This work was done while he was interning at OMRON SINIC X.

¹So Kuroki, Mai Nishimura, and Tadashi Kozuno are with OMRON SINIC X Corporation, 5-24-5, Hongo, Bunkyo-ku, Tokyo, Japan {mai.nishimura, tadashi.kozuno}@sinicx.com

the next-step actions of agents by incorporating local and global interactions among agents and a manipulation object. The prior demonstrations encompass a variety of scenarios with different object shapes (stick and block) and manipulation directions (pushing and rotating), which are both crucial for robot coordination in push manipulation tasks. These demonstrations are encoded into a single representative vector and stored in the coordination skill database as a pair of the vector and the original demonstration. Second, we develop a coordination skill matcher to measure similarities between two cooperative demonstrations in the learned skill representation space. Finally, we train the policy for the target task using a dataset augmented with the retrieved demonstrations. We demonstrate that our approach enables us to retrieve demonstrations that are beneficial for learning the task that requires similar coordination skills in both simulated and real-world scenarios.

We summarize our contributions as follows.

- We introduce a novel retrieve-and-learn framework for multi-agent systems which allows us to train a control policy using only a few target demonstrations.
- We introduce *Multi-Agent Coordination Skill Database*, a vector-based database constructed with a coordination skill encoder for storing and retrieving push manipulation behaviors, covering different object shapes and manipulation directions.
- We demonstrate that the skill database can provide reusable coordination skill representations not only for a simulated push manipulation task but also for a team of real-wheeled robots.

II. RELATED WORK

A. Multi-Agent Push Manipulations

Multi-agent push manipulation has received significant interest due to its substantial potential in object transportation, particularly in factory logistics and delivery services. In this task, it is crucial to acquire a complex interaction model that captures the relationships between agents and an object. Earlier research primarily concentrated on explicitly modeling interactions between robots and objects [5] or using rule-based methods to define robot movements [6]. More recent studies have shifted towards deep learning-based approaches, which leverage reinforcement learning or imitation learning to obtain cooperative interaction models [3], [4], [9], [10]. While these methods work well in simulated scenarios, they generally necessitate large datasets or long-time explorations, posing a challenge for real-world applications. In response, we introduce the Multi-Agent Coordination Skill Database based on interaction models distinctive to task-specific coordinated behaviors. The database enables the training of a multi-agent control policy using only a limited set of target demonstrations even in real-world situations.

B. Imitation Learning from Prior Demonstrations

Learning policies from experts have been extensively studied in various domains such as vehicle control [11] and multi-agent navigation [12]. Two major drawbacks of these

approaches are that the policy suffers from distribution shifts when adapted to unknown situations, and that the policy performs poorly when trained on sub-optimal demonstrations. Few-shot imitation learning is a promising approach for fine-tuning a policy, initially trained with a large-scale, multi-task dataset, to adapt to the target task using a small number of demonstrations [13], [14]. In contrast, the retrieval approach introduces a novel direction by using optimal demonstrations as a query to retrieve demonstrations relevant to the target task [7], [8]. Our work is built on top of their retrieval concept and is the first to expand this retrieve-and-learn paradigm to a multi-agent behavior database.

C. Constructing Multi-Task Prior Database

Following the success of large-scale pre-trained models built on extensive datasets, the robot learning community has developed open-source datasets [15]–[17]. These datasets feature demonstrations spanning multiple tasks and a range of robot types, gathered through methods like teleoperation [15], diverse learning approaches [16], and a combination of these methods inclusive of optimization techniques [17]. Despite these advancements, similar dataset development is not as prevalent in the multi-agent domain, particularly due to the inherent difficulty in controlling multi-agent robots associated with collision avoidance. To address these challenges, our work uses reinforcement learning in conjunction with the hybrid policy design incorporating reciprocal velocity obstacles (RVO) [18] to collect collision-free demonstrations of multiple agents at scale.

III. PROBLEM FORMULATION

Given a large task-agnostic prior dataset $\mathcal{D}_{\text{prior}}$ and a few demonstrations $\mathcal{D}_{\text{target}}$ collected from the target task, our main objective is to retrieve coordination skills from $\mathcal{D}_{\text{prior}}$ that facilitate downstream policy learning for the target task.

A. Multi-Agent Coordination Skill Database

$\mathcal{D}_{\text{prior}}$ consists of task-agnostic demonstrations of N mobile agents. Each demonstration, $\tau_i = (\mathcal{O}^1, \dots, \mathcal{O}^{T_i})$, is a sequence of variable length T_i , where $\mathcal{O}^t = \{(s_n^t, \mathbf{a}_n^t)\}_{n=1}^N$ is a set of N agents' state-action pairs for time step t . To effectively retrieve demonstrations from $\mathcal{D}_{\text{prior}}$ that are relevant to $\mathcal{D}_{\text{target}}$, we seek an abstract representation to measure the similarity between the two multi-agent demonstrations. We refer to this as *multi-agent coordination skill representation*, a compressed vector representation that is distinctive to the specific coordination behavior. For N agents' states s^t , we aim to learn a skill encoder \mathcal{E} that maps the state representation of multiple agents into a single representative vector $z \in \mathbb{R}^n$. Once the coordination skill representation space is learned on the prior dataset $\mathcal{D}_{\text{prior}}$, we construct a Multi-Agent Coordination Skill Database, a vector database consisting of the skill representation z associated with the original state-action pairs \mathcal{O} .

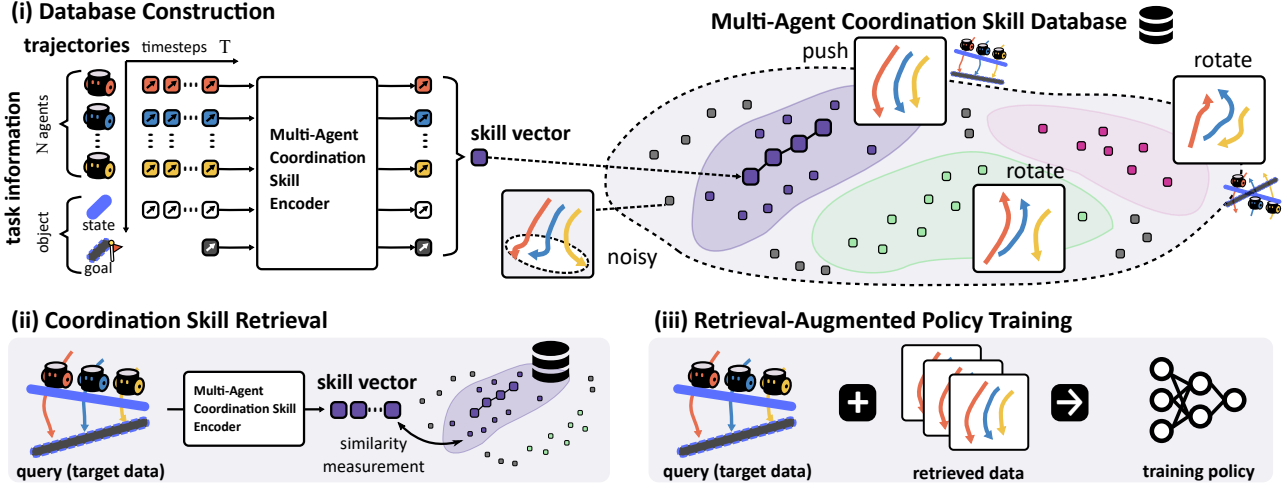


Fig. 2: Our retrieve-and-learn framework consists of three primary components: (a) constructing the Multi-Agent Coordination Skill Database on the basis of the prior experiences, (b) retrieving demonstrations using a few target demonstrations as queries, and (c) learning the multi-agent control policy using the retrieved data and target data.

B. Retrieval-Augmented Policy Training

We assume that the prior dataset $\mathcal{D}_{\text{prior}}$ and the target demonstrations $\mathcal{D}_{\text{target}}$ are composed as follows,

- $\mathcal{D}_{\text{prior}}$: The prior dataset encompasses a large-scale offline demonstration of diverse cooperative tasks. The dataset also includes noisy or sub-optimal demonstrations to mimic real-world scenarios. All demonstrations are task-agnostic, *i.e.*, each data is stored without any specific task annotations.
- $\mathcal{D}_{\text{target}}$: The target dataset encompasses a small amount of expert data, *i.e.*, all data is composed of well-coordinated demonstrations to complete the target task.

Given $\mathcal{D}_{\text{prior}}$ and $\mathcal{D}_{\text{target}}$, we aim to retrieve cooperative behaviors similar to those seen in $\mathcal{D}_{\text{target}}$ in the learned coordination skill space. Once we obtain the retrieved data \mathcal{D}_{ret} , we train a multi-agent control policy using $\mathcal{D}_{\text{target}}$ augmented with \mathcal{D}_{ret} . That is, the training data $\mathcal{D}_{\text{train}}$ is described as $\mathcal{D}_{\text{train}} = \mathcal{D}_{\text{target}} \cup \mathcal{D}_{\text{ret}}$.

C. Task Formulation

We consider the task of cooperative push manipulation where the objective is to navigate N mobile agents to push an object from the start to the goal while also interacting with each other. We assume a centralized policy [12] for controlling these agents and use imitation learning to train the policy. The objective of imitation learning is to learn a policy π that imitates the cooperative behavior of an expert π^* . Given offline demonstrations composed of state-action pairs \mathcal{O} , we aim to learn a cooperative policy $\pi : s^t \mapsto a^t$ that maps the multi-agent states to a set of their future actions. The state s^t is defined as $s^t = [\{r_n^t\}_{n=1}^N, o^t, g]$, where r_n^t denotes each agent's state and o^t is that of a manipulation object, respectively. We set the goal state g as a constant in a sequence. Each agent state r is defined as $r = [p, q_{\text{rot}}, w, h, v_{\text{prev}}]$, where p is the 2D position of the agent relative to the goal state coordinate, q_{rot} is rotation

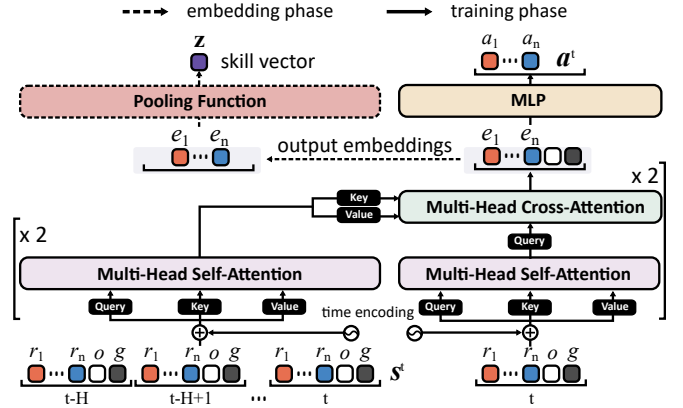


Fig. 3: In the training phase, our model takes past trajectory sequences as input tokens and outputs the future trajectories *i.e.*, actions of each input token. In the embedding phase, the model pools the multi-agent feature embeddings into a single representative vector.

to the world coordinate system represented by quaternion, w and h specify the width and height of the bounding box defined in the xy -plane in accordance with an agent, and v_{prev} is 2D velocity in the previous timestep. The object state o follows the same state components as the agent r . The next-step action a^t of each agent is defined as 2D velocity.

IV. PROPOSED METHOD

Figure 2 illustrates an overview of our proposed framework. The primary challenge of our retrieve-and-learn framework is to develop coordination skill representation that can discern similar or dissimilar coordinated behaviors of multiple agents. To address this, we propose a coordination skill encoder based on an attention mechanism. The encoder learns interactions between agents and a manipulation object from diverse prior demonstrations.

Furthermore, robot collisions, especially in the real world,

pose a significant challenge in tasks that require multi-agent coordination. We propose a controller that integrates a learning-based high-level policy with an optimization-based low-level policy. This design ensures safety by effectively avoiding robot collisions and makes our overall framework applicable to diverse environments.

A. Multi-Agent Coordination Skill Encoder

We introduce a Transformer-based coordination skill encoder, which learns to capture interactions among agents as well as interactions between agents and a manipulation object. The design of the coordination skill encoder is guided by the following key insights. First, the coordination occurs in both spatial and temporal domains. For example, if one agent maneuvers to the right side to contact an object, the other one must contact the left side of the object while predicting and interpreting the intention of the first agent. Second, the skill representation not only encompasses past trajectories but also predicts future trajectories. This enables the coordination skill to guide downstream policies, suggesting the appropriate action to take in future timesteps. To obtain such predictable skill representation space, we train the Transformer-based skill encoder as a prediction model of the agent’s future trajectories *i.e.*, actions.

Model Overview: Figure 3 illustrates an overview of our proposed model. We follow an encoder-decoder design of the standard Transformer architecture [19], where the encoder summarizes contextual information, while the decoder predicts the next state of the input tokens, taking into account the encoded context. In our network, the past context encoder is comprised of stacked multi-head self-attention layers that learn to attend to past trajectories across spatial and temporal domains. The future trajectory decoder is composed of multi-head self-attention layers and subsequent stacked multi-head cross-attention layers that integrate the past trajectory information and input tokens. Given past trajectory sequences with a temporal time window H , $\mathbf{X} = (\mathbf{s}^{t-H}, \mathbf{s}^{t-H+1}, \dots, \mathbf{s}^t)$, our model is trained to predict a set of next-step actions of the agents $\mathbf{Y} = \{\mathbf{a}_n^t\}_{n=1}^N$.

Attention Mechanism for Spatio-Temporal Interaction Encoding: As Transformers take sequential data as inputs, we first decompose trajectories of multiple agents, the manipulation object, and the goal state into individual tokens across agents, an object, a goal state, and time domains. The attention mechanism takes three types of inputs: a set of queries \mathbf{Q} , keys \mathbf{K} , and values \mathbf{V} . Formally, the attention function can be written as the mapping of a query and a set of key-value pairs to an output,

$$\text{Attn}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax} \left(\frac{\mathbf{Q}^\top \mathbf{K}}{\sqrt{d_k}} \right) \mathbf{V}, \quad (1)$$

where $\sqrt{d_k}$ is a scaling factor set in accordance with the length of the tokens. Each query, key, and value vector $\mathbf{q} = W_q \mathbf{h}_s$, $\mathbf{k} = W_k \mathbf{h}_t$, $\mathbf{v} = W_v \mathbf{h}_t$ are calculated from the input state sequence \mathbf{h} by a linear projection matrix W . We refer to the case of $\mathbf{h}_s = \mathbf{h}_t$ as self-attention and $\mathbf{h}_s \neq \mathbf{h}_t$ as cross-attention. To inform the model of the time sequence

order, we add time encoding to the input embeddings by sinusoidal encoding. In our setting, since the input tokens include positional information, the additional time encoding enables us to compute attention over spatial and temporal domains.

Training Loss: In the training phase, the obtained set of features $\mathbf{E} = \{e_1, \dots, e_N\}$ is decoded to a set of future actions by the multi-layer perception (MLP) layer. The skill encoder is trained to minimize the mean squared error (MSE) between ground truth actions a and predicted actions \hat{a} :

$$\text{MSE} = \frac{1}{T_i} \sum_{t=1}^{T_i} (a^t - \hat{a}^t)^2. \quad (2)$$

We train the coordination skill encoder using the whole $\mathcal{D}_{\text{prior}}$ involving diverse interactions by varying the number of agents. Our Transformer-based model accepts input tokens with variable lengths and is permutation-invariant in principle, which enables us to capture diverse interactions that are not limited to a specific number of agents.

Multi-Agent Feature Pooling: In the embedding phase, we pack the obtained features $\mathbf{E} = \{e_1, \dots, e_N\}$ into a single representative vector \mathbf{z} , *i.e.*, skill representation. To extract a representative feature invariant to both the specific number and order of agents, we employ a symmetric function. The function $\phi(\cdot)$ takes a set of vectors as input and outputs a single vector invariant to the input order,

$$\mathbf{z} = \phi(\{e_1, \dots, e_N\}). \quad (3)$$

For the symmetric function ϕ , we use an average pooling operation since it demonstrated better separation performance in the learned representation space compared with the max pooling operation. Finally, the Multi-Agent Coordination Skill Database is constructed by a collection of the skill vector \mathbf{z} associated with the original state-action pairs \mathcal{O} .

B. Data Retrieval from the Multi-Agent Coordination Skill Database

Given the query sampled from $\mathcal{D}_{\text{target}}$, our goal is to retrieve demonstrations that exhibit similar coordinated behavior to the query from the coordination skill database constructed from $\mathcal{D}_{\text{prior}}$. We first encode the query demonstrations into the learned skill representation space by the pre-trained skill encoder \mathcal{E} , and compare the similarity between the embedded query and collection of skill vectors in the database. The database stores sequences of skill vectors $\mathcal{Z}_j = (\mathbf{z}^1, \dots, \mathbf{z}^{T_j})$ with a variable length T_j embedded from diverse demonstrations in $\mathcal{D}_{\text{prior}}$. By encoding the query demonstrations through the skill encoder \mathcal{E} , we obtain a sequence of skill vectors with a variable length T'_j , $\mathcal{Z}'_j = (\mathbf{z}'^1, \dots, \mathbf{z}'^{T'_j})$. Note that these two sequences, \mathcal{Z}_i and \mathcal{Z}'_j , may have different lengths T_i and T'_j . To compare the two sequences of disparate lengths, we use dynamic time warping (DTW) [20], a renowned technique that can accommodate sequences of differing lengths to calculate the similarity

between temporal sequences.

$$DTW(\mathcal{Z}_i, \mathcal{Z}'_j) = \min_{\alpha \in A(\mathcal{Z}_i, \mathcal{Z}'_j)} \left(\sum_{(i,j) \in \alpha} m(\mathbf{z}^i, \mathbf{z}'^j) \right), \quad (4)$$

where m is the distance metric between the two elements, A contains all potential alignment paths between the two sequences, and α is the alignment path that pairs elements between sequences. In the experiment, we adopt a particular variant of DTW named *FastDTW* [21], which is recognized for its ability to approximate DTW in linear time, making it scalable for larger datasets. Assuming $\mathcal{D}_{\text{prior}}$ and $\mathcal{D}_{\text{target}}$ have L^p and L^t number of demonstrations, respectively, we compute the similarity for all combinations.

$$\begin{aligned} \text{FastDTW}(\mathcal{Z}_i, \mathcal{Z}'_j) \forall i \in \{0, 1, \dots, L^p - 1\} \\ \text{and } j \in \{0, 1, \dots, L^t - 1\} \end{aligned} \quad (5)$$

We utilize the cosine similarity as the distance metric m . We identify the top K similar demonstrations from $\mathcal{D}_{\text{prior}}$ to each target demonstration in $\mathcal{D}_{\text{target}}$, thereby retrieving a total of $L^t \times K$ demonstrations as \mathcal{D}_{ret} from $\mathcal{D}_{\text{prior}}$.

C. Hierarchical Policy for Safe Robot Control

Our ultimate goal is to deploy the cooperative policy π trained on the simulated dataset to real-world robots [22]. We employ a hierarchical policy design to ensure that our trained policy can navigate real-world mobile robots safely. The hierarchical planning policy is composed of i) a high-level policy that generates the agent’s velocities at sparse intervals and ii) a low-level control policy that adjusts the velocities outputted from the high-level policy to avoid potential collisions. We set the high-level control policy as our trained policy π and the low-level control policy as an off-the-shelf collision avoidance algorithm. We use the RVO algorithm with nonholonomic constraint [23]. That is, our policy can focus on the high-level planning of the cooperative behavior of each agent and can safely delegate collision avoidance to the external algorithm. This design enables us to make our overall framework versatile for typical wheeled robots with nonholonomic kinematic constraints.

V. EXPERIMENT

In this section, we experimentally evaluate our method using both simulation and real-world scenarios. Through our experiments, we aim to address the following questions:

- How is our multi-agent representation distinctly characterized for each skill? (V-D)
- Can our method retrieve coordinated behaviors similar to those in the prior dataset? (V-E)
- How does the retrieved dataset contribute to improving the success rate of the target task? (V-F and V-G)
- Can our method also enhance the control of multi-agent robots in real-world settings? (V-H)

A. Tasks and Dataset

Task Design: We assume N mobile robots are navigated to push an object toward a predefined goal state. To explore various coordination scenarios, we specifically focus on three different values of $N \in \{2, 3, 4\}$. For each of these setups, we collect demonstrations that involve four distinct tasks. These four tasks are carefully designed to encompass two different objects (**stick** or **block**) and manipulation difficulties (**easy** or **hard**). The stick and block objects are designed as 2D rectangles with dimensions of 0.24 m by 0.02 m and 0.24 m by 0.08 m, respectively. Stick objects can only be pushed stably from their longer side, while block objects can be pushed from either their longer or shorter side. These two different object designs result in diverse coordination scenarios. The difficulty of the task is defined by the relative distance d and the relative rotation θ between the initial and goal states. We refer to a **simple object rotating** task as **easy**, where the goal state is close to the initial state, ranging from 0 to 0.1 m, and θ ranges from -90° to 90° . A **combination of pushing and rotating** tasks is referred to as **hard**. As for hard tasks, the goal state is placed further from the initial state, with distances ranging from 0.1 to less than 0.2 m and θ ranges from -90° to 90° . At every setup, the initial position of the object is fixed, while the initial positions of the robots and the goal state are randomly sampled.

Dataset: We collect both the prior data $\mathcal{D}_{\text{prior}}$ and target data of $\mathcal{D}_{\text{target}}$ in the simulated environment implemented by Unity ML Agents [24]. We develop reinforcement learning-based policies based on soft actor-critic (SAC) [25] trained at various proficiency levels. To mimic real-world datasets involving sub-optimal demonstrations collected by non-experts, we prepared three distinct policies—classified as best, average, and worst—on the basis of their reward. This classification enabled the collection of a diverse dataset that reflects a spectrum of policy effectiveness. As for $\mathcal{D}_{\text{prior}}$, we amassed 2,500 demonstrations from each policy. For each N , this equates to $2,500 \times 3$ policies $\times 4$ tasks, totaling 30,000 demonstrations. As for $\mathcal{D}_{\text{target}}$, 250 demonstrations were collected from the best policy for each task.

B. Baselines

To the best of our knowledge, our work is the first to propose retrieving and learning coordination skill representations for multi-agent systems. Due to the lack of directly comparable existing methods, we have established two alternative baselines to assess our proposed method.

Agent-wise trajectory matching (A-TM): We utilize trajectory matching as a baseline to evaluate our retrieval method on the basis of multi-agent skill representation. By calculating the similarity of the robot’s xy -coordinate trajectories using *FastDTW*, we retrieve data from the $\mathcal{D}_{\text{prior}}$ that have trajectories similar to the target data’s trajectories.

Few-shot Imitation Learning (F-IL) [7], [8]: We compare our method with the adaptation method. We trained the multi-task policy from $\mathcal{D}_{\text{prior}}$ and fine-tuned it using $\mathcal{D}_{\text{target}}$.

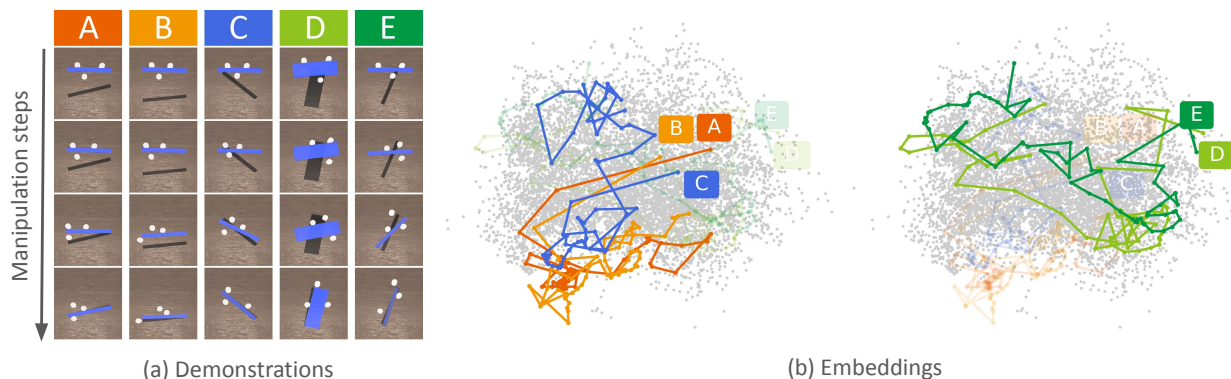


Fig. 4: (a) Visualization of five different trajectories ($N=3$): A, B, C, and E represent tasks of pushing a stick at a hard level, while D represents the task of pushing a block at a hard level. (b) Visualization of Multi-Agent Coordination representation. The left image highlights the representations of A, B, and C, while the right image highlights those of D and E. Each trajectory and its corresponding representation sequence are matched with both a unique color and an alphabet ID.

C. Training Details

In the implementation of our coordination skill encoder, we use two self-attention layers in the past context encoder and also two self-attention and cross-attention layers in the future trajectory decoder. We set the number of heads to 8 in all multi-head attention layers. The embedded dimensions of queries Q , keys K , and values V are all set to 256. We use an MLP layer with hidden dimensions of 512 and 256 to output the future trajectories. In the training phase, we use the Adam optimizer [26]. The learning rate and dropout rate are set to $1e-4$ and 0.1, respectively. We utilize the same model architecture for both the skill encoder and the multi-agent control policy. We trained the skill encoder with 10 epochs and all control policies with 50 epochs.

D. Visualization of the Multi-Agent Coordination Skill Representation Space

To accurately retrieve demonstrations that are similar to the target query, the skill representation should exhibit similarity in coordination behavior, while being dissimilar for unrelated tasks based on the object shapes and manipulation directions. We project each coordination skill representation into lower-dimensional space to view a property of the learned representation space. Figure 4 visualizes the first two dimensions of the embedding space projected by principal component analysis (PCA). Figure 4 (a) shows five example sequences of demonstrations, and Figure 4 (b) highlights their corresponding embeddings.

Similar Coordination Behaviors: In the left image of (b), A and B have the same representation space. This is because their behaviors are similar, as their goal states are alike. Observing the representations of A and C, they initially appear similar but later diverge. Trajectories A and C share similarities in coordination until the midpoint; beyond this, they diverge: A starts rotating the object counterclockwise, while C starts rotating it clockwise. Consequently, the representations of A and C initially mirror each other but later differ, reflecting their distinct behaviors.

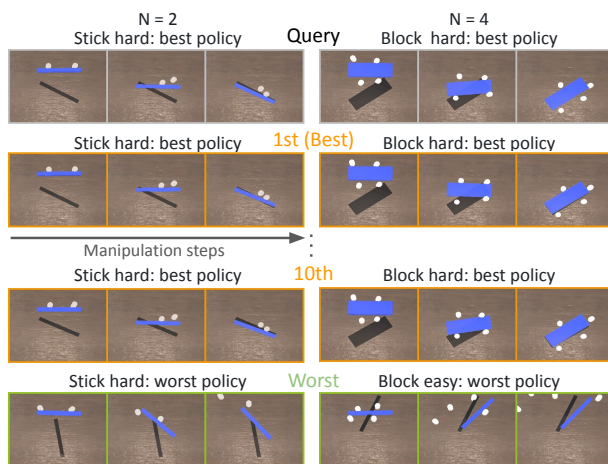


Fig. 5: The retrieved data is visualized with the two queries sampled from $\mathcal{D}_{\text{target}}$. For each demonstration, we display the manipulated object, the task difficulty, and the policy that generated the trajectory.

Dissimilar Coordination Behaviors: Comparing the left and right images, A, B, and C occupy distinctly separate spaces from D and E, with minimal overlap. This reflects the difference in trajectories between A, B, C and D, E.

Different Tasks: In the right images of (b), the representations of D and E have substantial overlap. Although D and E are pushing different objects, a stick and a block, respectively, the similarity in coordination is reflected in their overlapped representations.

E. Qualitative Results of Retrieved Data

In our retrieve-and-learn framework, all the retrieved demonstrations are ranked with the similarity metric, and the top-10 ranked demonstrations are used for downstream policy training. Figure 5 shows demonstrations ranked from the 1st (best) to 10th and worst retrieved by the query demonstration sampled from $\mathcal{D}_{\text{target}}$. The retrieved demonstrations ranked from 1st to 10th show similar movements to the query data. The demonstrations with the worst similarity possess

TABLE I: **Comparison with Baselines.** We evaluate the policies using 70 newly sampled test environments and calculate the success rate and standard error of the mean. Higher means and lower standard errors indicate better performance.

N	Task		Success Rate [%]		
	object	level	Ours	A-TM	F-IL
2	block	hard	41.4±5.9	41.4±5.9	20.0±4.8
	block	easy	57.1±6.0	55.7±6.0	15.7±4.4
	stick	hard	35.7±5.8	32.9±5.7	5.7±2.8
	stick	easy	67.1±5.7	68.6±5.6	12.9±4.0
3	block	hard	42.9±6.0	48.6±6.0	32.9±5.7
	block	easy	90.0±3.6	77.1±5.1	34.3±5.7
	stick	hard	41.4±5.9	28.6±5.4	34.3±5.7
	stick	easy	65.7±5.7	67.1±5.7	52.9±6.0
4	block	hard	58.6±5.9	55.7±6.0	38.6±5.9
	block	easy	88.6±3.8	78.6±4.9	55.7±6.0
	stick	hard	24.3±5.2	18.6±4.7	7.1±3.1
	stick	easy	70.0±5.5	52.9±6.0	57.1±6.0
ALL			56.9±1.7	52.1±1.7	30.6±1.6

goal states differing from the query data. Additionally, their trajectories are generated by the worst policy, which lacks effective coordination to complete tasks successfully. The result demonstrates that our method is capable of correctly ranking related movements higher and unrelated movements lower in relevance to the query data.

F. Evaluation of Downstream Policy Learning

Table I shows our retrieval-augmented policy training greatly outperforms F-IL [7], [8]. F-IL’s lower accuracy can likely be attributed to the multi-task policy, being adversely influenced by the noise data in $\mathcal{D}_{\text{prior}}$, and to the method’s inability to sufficiently fine-tune for the distribution shift with the limited data of $\mathcal{D}_{\text{target}}$. Our method also outperforms A-TM. As the number of robots N increases, the performance gap between our method and A-TM widens. This indicates that our approach is more effective in tasks requiring advanced robot coordination, particularly in scenarios with a larger number of robots or in more complex tasks.

G. Ablation Study

Efficacy of Retrieval-Augmented Policy Training: We conduct an ablation study to evaluate the effectiveness of our methods. To assess the role of the retrieval method, we train the policies using data solely from $\mathcal{D}_{\text{target}}$ (Target) and $\mathcal{D}_{\text{target}} \cup \mathcal{D}_{\text{prior}}$ (ALL). Table II shows that our method significantly outperforms these approaches. This suggests that our method, by appropriately retrieving data from $\mathcal{D}_{\text{prior}}$ on the basis of $\mathcal{D}_{\text{target}}$, is able to augment the training strategy effectively.

Distance Metric: We use the Euclidean distance as a similarity metric for searching and training the policy with the retrieved data. While the performance is fifty-fifty for each task, an overall comparison reveals that the proposed method, which employs cosine distance as the similarity metric, outperforms the policy using the Euclidean distance.

Balancing Loss Weights: Moreover, we modified the loss weights for behavior cloning on the basis of the similarity of the retrieved data. Given a list of similarity of

TABLE II: **Ablation study.** We evaluate the policies using 70 newly sampled test environments and calculate the success rate and standard error of the mean. Higher means and lower standard errors indicate better performance.

N	Task		Success Rate [%]				
	object	level	Ours	ALL	Target	Euclidean	Weight
2	block	hard	41.4±5.9	15.7±4.4	11.4±3.8	35.7±5.8	35.7 ± 5.8
	block	easy	57.1±6.0	22.9±5.1	17.1±4.5	47.1±6.0	44.3 ± 6.0
	stick	hard	35.7±5.8	17.1±4.5	18.6±4.7	38.6±5.9	34.3 ± 5.7
	stick	easy	67.1±5.7	32.9±5.7	7.1±3.1	70.0±5.5	65.7 ± 5.7
3	block	hard	42.9±6.0	14.3±4.2	12.9±4.0	57.1±6.0	45.7 ± 6.0
	block	easy	90.0±3.6	24.3±5.2	11.4±3.8	81.4±4.7	81.4 ± 4.7
	stick	hard	41.4±5.9	21.4±4.9	2.0±2.0	55.7±6.0	27.1 ± 5.4
	stick	easy	65.7±5.7	45.7±6.0	20.0±4.8	51.4±6.0	64.3 ± 5.8
4	block	hard	58.6±5.9	8.6±3.4	12.9±4.0	60.0±5.9	48.6 ± 6.0
	block	easy	88.6±3.8	51.4±6.0	21.4±4.9	85.7±4.2	77.1 ± 5.1
	stick	hard	24.3±5.2	1.4±1.4	4.3±2.4	25.7±5.3	18.6 ± 4.7
	stick	easy	70.0±5.5	25.7±5.3	8.6±3.4	55.7±6.0	75.7 ± 5.2
ALL			56.9±1.7	23.5±1.5	12.4±1.1	55.4±1.7	51.5 ± 1.7

the retrieved data l_1, l_2, \dots, l_K , the scaled loss weight is calculated as $S_i = \frac{l_{\min}}{l_i}$ for $i = 1, 2, \dots, K$, where S_i is the scaled loss weight corresponding to l_i . The result shows that the performance of this weighted behavior cloning is declining. This may be because the weighting has led to an undervaluation of beneficial trajectories during training.

H. Real Robot Experiments

To validate the efficacy of our method in the real world, we use demonstrations of real wheeled robots for querying the prior dataset $\mathcal{D}_{\text{prior}}$ constructed in simulated environments. We then train the policy using a few real-robot data augmented with the retrieved simulation demonstrations. We employ the custom-designed swarm robot platform [22], where the wheeled microrobots communicate with the host computer using 2.4 GHz wireless communication. The positions of all robots and an object are tracked in real time by a high-speed digital light processing (DLP) structured light projector system. We can control the positions of the robots by sending 2D coordinates of their future locations from the host computer.

We collect three demonstrations that navigate robots toward a given goal state by using a hand-gesture system with Leap Motion, where the hands’ movement is translated into the trajectories of multiple robots. Using the real-robot demonstrations as queries, we retrieve 300 simulated demonstrations from $\mathcal{D}_{\text{prior}}$ for each query. Figure 6 shows four randomly sampled retrieval demonstrations in response to a query of the real-robot trajectory. The result demonstrates that our method can accurately rank coordinated movements similar to the query higher. Figure 7 compares our policy trained using the real and retrieved demonstrations and the one trained using only the real-robot demonstrations. The results clearly demonstrate that the policy trained by our retrieve-and-learn framework successfully pushes an object closer to the goal state, while the policy trained using only a few real demonstrations fails to complete the task.

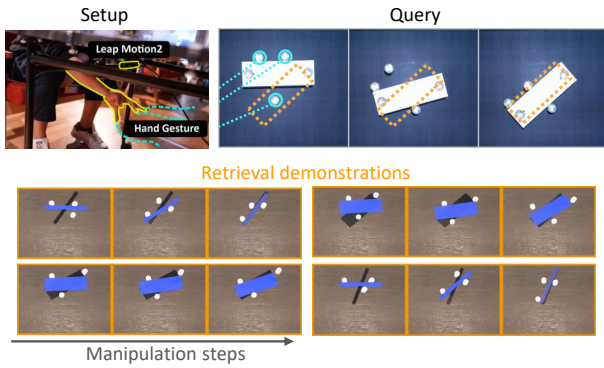


Fig. 6: Top Left: Hand gesture system used for collecting real demonstrations. Top Right: Example of a query retrieved using the system. Bottom: Four randomly selected examples from 300 demonstrations retrieved by the query.




Fig. 7: Comparison highlighting the differences between our retrieve-and-learn policy and one that was trained using only real-world target trajectories.

VI. CONCLUSION AND FUTURE WORK

We introduced a novel concept named Multi-Agent Coordination Skill Database, an essential skill representation repository for push manipulation distinctive to object shapes and its manipulating directions. Extensive evaluations in the simulated environments clearly demonstrate the effectiveness of our retrieve-and-learn approach on the basis of the skill database both qualitatively and quantitatively. There are two future directions. In real-world deployments, we have observed failure cases in scenarios involving complex collisions between multiple robots and a manipulation object. To address this issue, we believe that identifying the parameters through real-to-sim approaches or incorporating domain randomization techniques would be a promising direction. In addition, we aim to extend our framework's scope, examining its scalability for diverse push manipulation setups and its generalizability to other multi-agent tasks such as navigation and mobile-arm manipulation.

Acknowledgement: This work was supported by JST AIP Acceleration Research JPMJCR23U2, Japan.

REFERENCES

- [1] C. Yang, G. N. Sue, Z. Li, L. Yang, H. Shen, Y. Chi, A. Rai, J. Zeng, and K. Sreenath, "Collaborative navigation and manipulation of a cable-towed load by multiple quadrupedal robots," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 10 041–10 048, 2022.
- [2] H. M. Le, Y. Yue, P. Carr, and P. Lucey, "Coordinated multi-agent imitation learning," in *Proc. ICML*, 2017, pp. 1995–2003.
- [3] K. Shibata, T. Jimbo, and T. Matsubara, "Deep reinforcement learning of event-triggered communication and control for multi-agent cooperative transport," in *Proc. ICRA*, 2021, pp. 8671–8677.
- [4] S. Kuroki, T. Matsushima, J. Arima, H. Furuta, Y. Matsuo, S. S. Gu, and Y. Tang, "Collective intelligence for 2d push manipulations with mobile robots," *IEEE Robotics and Automation Letters*, vol. 8, no. 5, pp. 2820–2827, 2023.
- [5] Z. Wang and V. Kumar, "Object closure and manipulation by multiple cooperating mobile robots," in *Proc. ICRA*, vol. 1, 2002, pp. 394–399.
- [6] J. Chen, M. Gauci, W. Li, A. Kolling, and R. Groß, "Occlusion-based cooperative transport with a swarm of miniature mobile robots," *T-RO*, vol. 31, no. 2, pp. 307–321, 2015.
- [7] S. Nasiriany, T. Gao, A. Mandlekar, and Y. Zhu, "Learning and retrieval from prior data for skill-based imitation learning," in *Proc. CoRL*, 2023, pp. 2181–2204.
- [8] M. Du, S. Nair, D. Sadigh, and C. Finn, "Behavior retrieval: Few-shot imitation learning by querying unlabeled datasets," in *Proc. RSS*, 2023.
- [9] O. Nachum, M. Ahn, H. Ponte, S. Gu, and V. Kumar, "Multi-agent manipulation via locomotion using hierarchical sim2real," in *Proc. CoRL*, 2019.
- [10] J. Wu, X. Sun, A. Zeng, S. Song, S. Rusinkiewicz, and T. Funkhouser, "Spatial intention maps for multi-agent mobile manipulation," in *Proc. ICRA*, 2021, pp. 8749–8756.
- [11] R. P. Bhattacharyya, D. J. Phillips, B. Wulfe, J. Morton, A. Kuefler, and M. J. Kochenderfer, "Multi-agent imitation learning for driving simulation," in *Proc. IROS*, 2018, pp. 1534–1539.
- [12] K. Renz, K. Chitta, O.-B. Mercea, A. Koepke, Z. Akata, and A. Geiger, "Plant: Explainable planning transformers via object-level representations," in *Proc. CoRL*, 2022, pp. 459–470.
- [13] Z. Mandi, F. Liu, K. Lee, and P. Abbeel, "Towards more generalizable one-shot visual imitation learning," in *Proc. ICRA*, 2022, pp. 2434–2444.
- [14] H.-S. Fang, H. Fang, Z. Tang, J. Liu, J. Wang, H. Zhu, and C. Lu, "Rh20: A robotic dataset for learning diverse skills in one-shot," in *RSS 2023 Workshop on Learning for Task and Motion Planning*, 2023.
- [15] A. Mandlekar, Y. Zhu, A. Garg, J. Booher, M. Spero, A. Tung, J. Gao, J. Emmons, A. Gupta, E. Orbay, *et al.*, "Roboturk: A crowdsourcing platform for robotic skill learning through imitation," in *Proc. CoRL*, 2018, pp. 879–893.
- [16] J. Fu, A. Kumar, O. Nachum, G. Tucker, and S. Levine, "D4rl: Datasets for deep data-driven reinforcement learning," *CoRR*, 2020.
- [17] Q. Vuong, S. Levine, H. R. Walke, K. Pertsch, A. Singh, R. Doshi, C. Xu, J. Luo, L. Tan, D. Shah, *et al.*, "Open x-embodiment: Robotic learning datasets and rt-x models," in *Towards Generalist Robots: Learning Paradigms for Scalable Skill Acquisition@ CoRL*, 2023.
- [18] J. Van den Berg, M. Lin, and D. Manocha, "Reciprocal velocity obstacles for real-time multi-agent navigation," in *Proc. ICRA*, 2008, pp. 1928–1935.
- [19] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. NeurIPS*, vol. 30, 2017.
- [20] M. Müller, "Dynamic time warping," *Information retrieval for music and motion*, pp. 69–84, 2007.
- [21] S. Salvador and P. Chan, "Toward accurate dynamic time warping in linear time and space," *Intelligent Data Analysis*, vol. 11, no. 5, pp. 561–580, 2007.
- [22] S. Ichihashi, S. Kuroki, M. Nishimura, K. Kasaura, T. Hiraki, K. Tanaka, and S. Yoshida, "Swarm body: Embodied swarm robots," in *Proc. CHI*, 2024.
- [23] J. Snape, J. Van Den Berg, S. J. Guy, and D. Manocha, "Smooth and collision-free navigation for multiple robots under differential-drive constraints," in *Proc. IROS*, 2010, pp. 4584–4589.
- [24] A. Juliani, V.-P. Berges, E. Teng, A. Cohen, J. Harper, C. Elion, C. Goy, Y. Gao, H. Henry, M. Mattar, and D. Lange, "Unity: A general platform for intelligent agents," *CoRR*, 2020.
- [25] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *Proc. ICML*, 2018, pp. 1861–1870.
- [26] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. ICLR*, 2015.