# FROM SOLO TO SYMPHONY: ORCHESTRATING MULTI-AGENT COLLABORATION WITH SINGLE-AGENT DEMOS

**Anonymous authors**Paper under double-blind review

# **ABSTRACT**

Training a team of agents from scratch in multi-agent reinforcement learning (MARL) is highly inefficient, much like asking beginners to play a symphony together without first practicing solo. Existing methods, such as offline or transferable MARL, can ease this burden, but they still rely on costly multi-agent data, which often becomes the bottleneck. In contrast, solo experiences are far easier to obtain in many important scenarios, e.g., collaborative coding, household cooperation, and search-and-rescue. To unlock their potential, we propose Solo-to-Collaborative RL (SoCo), a framework that transfers solo knowledge into cooperative learning. SoCo first pretrains a shared solo policy from solo demonstrations, then adapts it for cooperation during multi-agent training through a policy fusion mechanism that combines an MoE-like gating selector and an action editor. Experiments across diverse cooperative tasks show that SoCo significantly boosts the training efficiency and performance of backbone algorithms. These results demonstrate that solo demonstrations provide a scalable and effective complement to multi-agent data, making cooperative learning more practical and broadly applicable.

# 1 Introduction

Multi-agent reinforcement learning (MARL) has emerged as a core paradigm for sequential decision making in environments that require coordination (Shoham & Leyton-Brown, 2008; Lowe et al., 2017; Gronauer & Diepold, 2022). By interacting with the environment and receiving feedback, MARL enables agents to learn cooperative policies, providing a principled framework for solving complex decision-making problems such as autonomous driving (Zhang et al., 2024), large-scale network optimization (Stepanov et al., 2024), and collaborative robotics (Tang et al., 2025).

However, compared to single-agent RL, MARL faces inherent challenges (Busoniu et al., 2008; Hernandez-Leal et al., 2019), including dimensionality explosion, coordination difficulty, and environmental non-stationarity. As a result, training joint policies from scratch is often inefficient, much like asking novices to rehearse a symphony without prior practice: difficult, time-consuming, and unlikely to yield good results. This inefficiency poses a major obstacle to applying MARL effectively in practice.

To address these challenges, a growing line of research has explored offline MARL (Pan et al., 2022; Shao et al., 2023; Liu et al., 2023; Liu et al., 2024b) and offline-to-online fine-tuning (Zhong et al., 2025). These methods learn from pre-collected task-specific cooperative trajectories to avoid costly exploration, and refine pretrained policies with limited online rollouts when interaction is allowed. More recent studies have attempted to relax the data assumption by leveraging multi-task cooperative datasets (Zhang et al., 2023a; Chen et al., 2024; Liu et al., 2025) or even non-cooperative multi-agent datasets (Wang et al., 2023; Yu et al., 2025). These efforts broaden the scope of usable data and represent important progress, but they remain fundamentally tied to multi-agent trajectories.

Actually, in many cooperative problems, there often exists a corresponding solo version whose demonstrations are much easier to obtain and learn from. For example, in collaborative coding (Dong et al., 2025b) a single coder writes a short piece of code, in household cooperation (Kannan et al., 2024) a single robot performs an individual chore, and in search-and-rescue (Cao et al., 2024)

a single drone searches for a target. Although such demonstrations deviate from the target cooperative setting, they are far from useless. Such as in orchestral performance, it is more effective to let each novice player first master the basics of solo play before attempting a full ensemble. Yet, the potential of solo demonstrations to accelerate MARL remains underexplored. This gap motivates an important but underexplored question:

# Can solo demonstrations be effectively leveraged to accelerate the collaborative MARL?

An affirmative answer to this hypothesis will validate solo data as a scalable and cost-effective resource. This will be instrumental in fostering efficient learning in settings where cooperative data are limited but solo demonstrations are plentiful (Kannan et al., 2024; Cao et al., 2024; Dong et al., 2025b), consequently making MARL a more viable solution for practical applications.

However, addressing this problem is non-trivial and involves two major challenges. The first is *observation mismatch*: differences in observation dimensionality hinder the direct transfer of solo demonstrations to multi-agent training (Hu et al., 2021; Zhang et al., 2023a; Liu et al., 2025). In some cases, a single local observation may even correspond to multiple distinct solo views, creating ambiguity for policy reuse. The second is *domain shift*: unlike multi-agent data, whether joint or agent-specific, that inherently encode cooperation (Wang et al., 2023; Yu et al., 2025), solo data contain no such information. In addition, discrepancies in environment dynamics between solo and cooperative settings (e.g., individual attributes and observation noise) further exacerbate the gap. These challenges hinder direct policy transfer, highlighting the need to distill knowledge from solo demonstrations and integrate it into cooperative learning.

To tackle these challenges, we propose <u>Solo-to-Collaborative</u> RL (SoCo) framework, which transfers knowledge from solo demonstrations to cooperative MARL. SoCo first pretrains a shared solo policy from solo demonstrations via imitation learning, providing a common skill prior for all agents. During cooperative training, local observations are decomposed into solo views aligned with the demonstrations, allowing the reuse of the solo policy to obtain candidate actions. Then, a policy fusion module selects and refines these actions for each agent, adapting them to the cooperative setting and mitigating domain shift. Specifically, inspired by MoE (Cai et al., 2025) and action fusion in single-agent RL (Dong et al., 2025a), a learnable gating selector chooses the most suitable candidate, while an action editor refines it for effective cooperation. This design not only tackles the challenge of solo-to-cooperative transfer but also provides flexibility for task-specific customization.

We validate SoCo across diverse cooperative benchmarks, and the results show that it markedly improves training efficiency while achieving competitive or superior performance. These findings highlight the potential of solo demonstrations as a scalable resource for cooperative MARL.

Our main contributions are summarized as follows:

- We investigate an important yet underexplored problem of leveraging solo demonstrations
  to benefit cooperative MARL, and show that such data, though lacking explicit cooperative
  information, can substantially accelerate multi-agent training.
- We develop Solo-to-Collaborative RL (SoCo), a framework for solo-to-cooperative transfer. It decomposes local observations to reuse a pretrained shared solo policy, and employs a policy fusion module trained from cooperative interactions that combines a gating selector for choosing solo actions with an action editor for refining them, enabling more efficient cooperation.
- We validate SoCo on cooperative benchmarks with diverse characteristics and difficulty, showing that it effectively addresses observation ambiguity and domain shift, boosts training efficiency, and achieves competitive or superior performance, highlighting the potential of solo demonstrations as a scalable resource for MARL.

# 2 Preliminary

#### 2.1 Multi-Agent Reinforcement Learning

We model multi-agent reinforcement learning (MARL) within the decentralized partially observable Markov decision process (Dec-POMDP) framework (Oliehoek et al., 2016). Formally, a Dec-

POMDP is defined as  $\mathcal{M} = \langle \mathcal{N}, \mathcal{S}, \mathcal{A}, \mathcal{O}, P, R, \gamma \rangle$ , where  $\mathcal{N} = \{1, \dots, N\}$  denotes the set of agents,  $\mathcal{S}$  is the global state space, and  $\mathcal{A}$  and  $\mathcal{O}$  are the joint action and observation spaces, each formed from the agents' local action  $\{\mathcal{A}_i\}_{i=1}^N$  and observation spaces  $\{\mathcal{O}_i\}_{i=1}^N$ . At each time step, every agent receives a local observation generated from the current global state. Based on the joint action  $\mathbf{a} = (a_1, \dots, a_N)$ , the environment transitions to the next state according to P, and a shared reward  $R(s, \mathbf{a})$  is returned. The goal of MARL is to learn a joint policy  $\mathbf{\pi} = (\pi_1, \dots, \pi_N)$  that maximizes the expected discounted return:  $J(\mathbf{\pi}) = \mathbb{E}_{\mathbf{\pi}}[\sum_{t=0}^{\infty} \gamma^t R(s_t, \mathbf{a}_t)]$ . The solo case naturally arises when  $|\mathcal{N}| = 1$ .

#### 2.1.1 CTDE PARADIGM AND DETERMINISTIC POLICY GRADIENT METHOD

Centralized Training with Decentralized Execution (CTDE) (Oliehoek et al., 2008; Amato, 2024; Li et al., 2025b) is a widely adopted paradigm in cooperative MARL. In CTDE, each agent executes its policy in a decentralized manner, relying only on its own local observation during interaction with the environment. During training, however, the learning process can leverage additional global information (e.g., global states or joint actions) through centralized critics. This design improves training stability and coordination, while keeping execution scalable and realistic.

A representative CTDE algorithm is MADDPG (Lowe et al., 2017). It extends the deterministic policy gradient (DPG) framework to multi-agent settings by introducing a centralized critic for each agent, while keeping actors decentralized. Formally, let agent i have policy  $\pi_i(o_i; \theta_i)$  parameterized by  $\theta_i$ , and the deterministic policy gradient for agent i is:

$$\nabla_{\theta_i} J(\pi_i) \approx \mathbb{E}_{s, \mathbf{a} \sim \mathcal{D}} \left[ \nabla_{\theta_i} \pi_i(o_i) \nabla_{a_i} Q_i(s, \mathbf{a}; \psi^i) \Big|_{a_i = \pi_i(o_i)} \right],$$

where  $Q_i(s, \boldsymbol{a}; \psi^i)$  denotes the centralized critic for agent i, parameterized by  $\psi^i$ . In practice, however, the critics often share a single parameter set  $\psi$  across agents, whereas each agent maintains its own policy network.

Building on MADDPG, MATD3 (Ackermann et al., 2019) incorporates the improvements of TD3 (Fujimoto et al., 2018), including twin critics, target smoothing, and delayed policy updates. HATD3 (Zhong et al., 2024) further extends MATD3 by introducing a heterogeneous sequential optimization. In this paper, we focus primarily on the DPG family under the CTDE paradigm, as represented by MATD3 and HATD3. Nevertheless, the proposed framework is, in principle, extendable to stochastic policy methods, such as MAPPO (Yu et al., 2022) or HASAC (Liu et al., 2024a).

#### 3 SOLO-TO-COLLABORATIVE RL

To bridge the gap between solo demonstrations and multi-agent cooperation, and to tackle observation mismatch and domain shift, we propose the Solo-to-Collaborative RL (SoCo) framework. SoCo first learns a shared solo policy from solo demonstrations. Then, during cooperative training, local observations are decomposed into solo views, allowing the reuse of the solo policy to obtain candidate actions. Finally, a per-agent policy-fusion module selects the most appropriate candidate policy and refines it for each agent, adapting it to the cooperative setting and mitigating domain shift. In the following, we present each component in detail. The full algorithm is presented in Algorithm 1 in Appendix C.

# 3.1 Solo Policy Extraction

A solo policy is learned by imitation from the demonstration dataset  $\mathcal{D}_s$  and shared by all the agents. For simplicity, our implementation uses standard behavior cloning, minimizing the mean-squared error between the policy's action and the action recorded in  $\mathcal{D}_s$  to obtain a deterministic behavioral policy:

$$\min_{av} \mathbb{E}_{(o,a)\sim\mathcal{D}_s} \left\| \beta_w(o) - a \right\|_2^2. \tag{1}$$

This design choice is flexible rather than mandatory: one could instead adopt a stochastic imitation model that learns  $\beta_w(a \mid o)$ , for example, parameterizations based on VAE (Kingma & Welling, 2013), diffusion models (Yang et al., 2023), or flow matching (Lipman et al., 2022), by simply

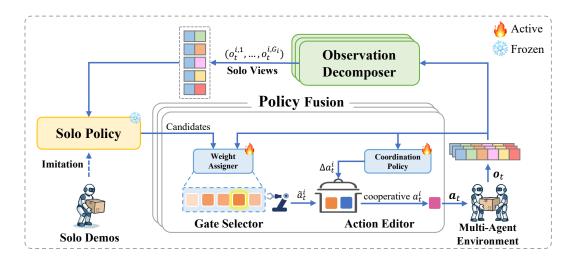


Figure 1: SoCo framework. A shared solo policy is pretrained from demonstrations and kept frozen, then reused through observation decomposition during cooperative training. Coordination ability is injected by the Policy Fusion module, where the Gating Selector selects suitable solo actions and the Action Editor fine-tunes them to mitigate domain shift.

switching to a likelihood-based objective without altering the subsequent components of the framework. Finally, the solo policy is shared across agents, and its parameters are frozen during the cooperative learning phase.

# 3.2 OBSERVATION DECOMPOSITION

Given the settings of the cooperative tasks and their corresponding solo tasks considered in this paper, we make a reasonable assumption that the observations in these cooperative tasks are well-defined, structured, and decomposable. Specifically, each observation consists of own features (e.g., velocity, position) and stacked features of all other entities (e.g., teammates or target states). The observation space of the corresponding solo tasks can then be constructed from these feature units (e.g., controlling one HalfCheetah vs. multiple coupled HalfCheetahs).

Hence, following prior works (Liu et al., 2020; Wu et al., 2024; Liu et al., 2025), we introduce a rule-based observation decomposer. Concretely, we decompose the observation of the i-th agent at time step t, denoted as  $o_t^i$ , into the self-related component  $o_t^{i,0}$  and the entity-related components  $\{o_t^{i,k}\}_{k=1}^{K_i}$ , where  $K_i$  denotes the total number of entities observable by agent i. When deployed in cooperative environments, depending on the specific task, we may reassemble the decomposed feature units into  $G_i$  valid solo views  $\{\tilde{o}_t^{i,k}\}_{k=0}^{G_i}$  for agent i by concatenation, zero-padding, and so on, thereby addressing the issue of inconsistent observation spaces.

#### 3.3 POLICY FUSION

In the cooperative training phase, the pretrained solo policy cannot be directly transferred. The main obstacles are twofold: (i) a single local observation may map to multiple solo views, producing several candidate actions (e.g., toward different targets) that must be disambiguated; and (ii) domain shift between solo and multi-agent settings often degrades performance, necessitating fine-tuning for effective adaptation.

Therefore, inspired by Mixture-of-Experts (MoE) (Cai et al., 2025) and action fusion techniques in single-agent RL (Dong et al., 2025a), we propose a novel learnable policy fusion module. Notably, our design operates at the agent level and is trained directly on multi-agent samples with standard MARL optimization, thereby injecting cooperative adaptability into solo policies. Within this module, each agent employs a *Gating Selector* to resolve ambiguity by selecting suitable solo actions, and an *Action Editor* to fine-tune the chosen action for coordination, together enabling effective solo-to-cooperative transfer.

#### 3.3.1 GATING SELECTOR

As discussed in Section 3.2, the local observation  $o_t^i$  of the *i*-th agent corresponds to  $G_i$  solo views  $\{\tilde{o}_t^{i,k}\}_{k=0}^{G_i}$ . By applying the solo policy, these yield  $G_i$  candidate actions:

$$a_t^{i,k} = \beta(\tilde{o}_t^{i,k}), \quad k = 1, \dots, G_i.$$

However, due to the solo-to-cooperative gap, not all candidate actions are suitable for the current cooperative context, and some may even conflict with each other. To resolve this, SoCo equips each agent with a weight assigner  $g_{\varphi}^i:\mathcal{O}_i\to\mathbb{R}^{G_i}$  that, conditioned on the current local observation  $o_t^i$ , evaluates the candidate solo actions and assigns weights to them, thereby selecting the most appropriate one for coordination.

To enable learnable action selection, we adopt the Gumbel–Softmax reparameterization (Jang et al., 2017) with the straight-through estimator. The gating weights  $g_{\varphi}^{i}(o_{t}^{i})$  define a categorical distribution, from which a one-hot action is drawn: the forward pass takes the most probable action, while the backward pass propagates gradients through the soft sample. The resulting action for agent i is:

$$\tilde{a}_t^i = \langle \text{GumbelSoftmax}(g_{\omega}^i(o_t^i)), \, \boldsymbol{a}_t^i \rangle,$$
 (3)

where  $a_t^i = (a_t^{i,1}, \dots, a_t^{i,G_i})$  is the set of candidate actions derived from the solo policy.

Moreover, this module is designed to be both general and flexible, allowing adaptation to different scenarios. For instance, the gating function may be rule-based instead of learned, and in the special case of  $G_i = 1$ , the selector can be omitted entirely.

#### 3.3.2 ACTION EDITOR

To leverage the prior knowledge in solo actions while overcoming transfer difficulties from domain shift, we design an action editor that injects cooperative information through residual corrections. Specifically, we introduce a coordination policy  $\pi_{\theta}: \mathcal{O}_i \to \mathcal{A}_i$  that produces a raw residual adjustment to the solo action. To keep this correction bounded and scale-invariant while avoiding gradient saturation, we squash the policy output with  $f_L(x) = L \tanh(x/L)$ . Given the current local observation  $o_i^t$ , the adjustment is:

$$\Delta a_t^i = \begin{cases} L \cdot \tanh\left(\frac{\pi_\theta(o_t^i)}{L}\right) & \text{if } L > 0\\ \mathbf{0} & \text{if } L = 0 \end{cases}$$
(4)

where L is a hyperparameter that controls the strength of the correction. By tuning L, the framework can trade off between leveraging solo priors and adapting to multi-agent dynamics.

Then, the final cooperative action is defined as:

$$a_t^i = \operatorname{Clip}(\tilde{a}_t^i + \Delta a_t^i) \tag{5}$$

where  $\tilde{a}_t^i$  denotes the solo action selected by the gating selector, and  $\mathrm{Clip}(\cdot)$  is a clipping operator to prevent action overflow. In our implementation, we adopt a  $\mathrm{tanh}$ -based operator, but the design is modular and allows substituting other operators depending on the task.

# 3.4 COLLABORATIVE POLICY OPTIMIZATION

For notational simplicity, we denote the fused policy as  $\Pi_{\phi}$ , where  $\phi = \{\varphi, \theta\}$  collects the learnable parameters of the gating selector and action editor. Since SoCo is fully decoupled from the backbone algorithm, this policy can be optimized with any MARL method; here we instantiate it with MATD3 (Ackermann et al., 2019) for concreteness.

Each agent i maintains two shared centralized critics  $Q_{\psi_1}, Q_{\psi_2}$  and an individual actor  $\Pi_{\phi_i}$ . Given a batch of transition  $\mathcal{B} = \{(s_t, \boldsymbol{o}_t, \boldsymbol{a}_t, r_t, s_{t+1}, \boldsymbol{o}_{t+1})\}$ , the critic loss is:

$$\mathcal{L}(\psi_j) = \mathbb{E}_{(s_t, \boldsymbol{o}_t, \boldsymbol{a}_t, r_t, \boldsymbol{o}_{t+1}) \sim \mathcal{B}} \left[ \left( Q_{\psi_j}(s_t, \boldsymbol{a}_t) - y_t \right)^2 \right], \quad j = 1, 2,$$
 (6)

where the target  $y_t$  is defined as

$$y_t = r_t + \gamma \min_{k=1,2} Q_{\bar{\psi}_k} (s_{t+1}, \mathbf{a}'_{t+1}), \quad (\mathbf{a}'_{t+1})_i = \Pi_{\bar{\phi}_i} (o^i_{t+1}) + \epsilon, \tag{7}$$

with  $\{\bar{\psi}_k\}_{k=1}^2$  and  $\bar{\phi}_i$  denoting target networks and  $\epsilon$  being clipped Gaussian noise for policy smoothing. The actors are optimized by maximizing the Q-value estimated by the first critic:

$$\mathcal{L}(\phi_i) = -\mathbb{E}_{(s_t, \boldsymbol{o_t}) \sim \mathcal{B}} \left[ Q_{\psi_1} \left( s_t, \Pi_{\phi_i}(o_t^i) \right) \right]. \tag{8}$$

For algorithms with stochastic policies, it suffices to compute the log-probability of the fused action according to Eq. (5) and substitute it into the policy loss.

In this way, SoCo provides a plug-and-play bridge between solo demonstrations and cooperative learning, turning single-agent demonstrations into a scalable and effective complement to multiagent data, making cooperative learning more practical and broadly applicable.

# 4 EXPERIMENTS

To evaluate the proposed SoCo framework, we conduct experiments on a variety of cooperative tasks. Our goals are to investigate the following questions: (i) Can SoCo improve the sample efficiency of multi-agent algorithms? (ii) Can SoCo enhance the ultimate performance of multi-agent algorithms? (iii) How do the individual components of SoCo contribute to its effectiveness?

#### 4.1 SETUP

Environments and Tasks. Following prior works (Sun et al., 2023; Kontogiannis et al., 2025; Zeng et al., 2025), our experiments cover nine tasks from four representative cooperative scenarios: (i) *Spread* (Lowe et al., 2017; Terry et al., 2021), where agents must cover distinct landmarks. This setting introduces target ambiguity but involves little domain shift. (ii) *LongSwimmer* (Peng et al., 2021; de Lazcano et al., 2024), where a multi-segment worm must swim forward, with each agent controlling two consecutive joints. These tasks do not involve target ambiguity but introduce domain shift due to altered dynamics. (iii) *MultiHalfCheetah* (Peng et al., 2021; de Lazcano et al., 2024), where multiple HalfCheetahs are connected in a chain and must run forward together. These tasks avoid target ambiguity but involve noticeable domain shift and present a non-trivial control challenge. (iv) *MultiWalker* (Gupta et al., 2017; Terry et al., 2021), where multiple bipedal robots jointly carry a package forward. These tasks avoid target ambiguity, but are inherently very difficult, with severe domain shift on top of coordination challenges.

Considering the characteristics and difficulty of these environments, we evaluate tasks with 3, 4, and 5 agents in *Spread* and *LongSwimmer*, 2 and 3 agents in *MultiHalfCheetah*, and 2 agents in *MultiWalker*. Details on these environments and tasks are provided in Appendix D.1.

**Data Collection.** To collect solo demonstration data, we first train policies with TD3 (Fujimoto et al., 2018) on the corresponding solo tasks until convergence, and then record 1M transition samples. The corresponding solo tasks are: (i) a single agent navigating to one landmark, (ii) a 3-segment worm with 2 joints swimming forward, (iii) a HalfCheetah with altered attributes running forward, and (iv) a single bipedal robot carrying the long package forward. When transferred to multi-agent settings, these lead to (i) goal ambiguity, (ii) domain shift, (iii) notable domain shift and cooperative difficulty, and (iv) severe domain shift coupled with substantial cooperative difficulty. More details can be found in Appendix D.2.

**Baselines.** We adopt two representative DPG algorithms, MATD3 (Ackermann et al., 2019) and HATD3 (Zhong et al., 2024), as our baselines. Their implementations are taken from the HARL codebase (Zhong et al., 2024), on top of which we also build our SoCo variants.

**Experiment Setup.** SoCo first undergoes imitation learning to obtain the solo policy (5k steps for *Spread*, 100k steps for the others). During cooperative learning, all algorithms perform 10k random interaction steps for warm-up, followed by policy optimization with the same number of environment steps (2M steps for *LongSwimmer* and *MultiHalfCheetah*, 5M steps for the others).

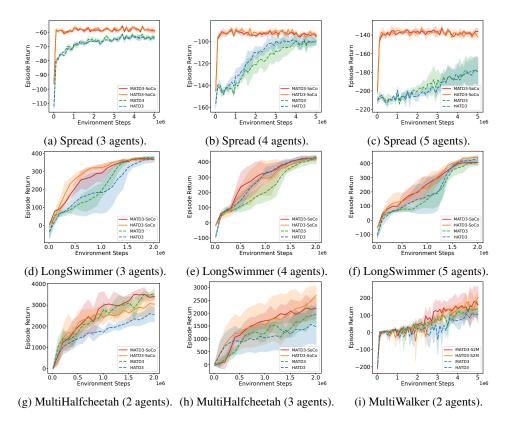


Figure 2: Training curves on nine tasks. Results are averaged over three random seeds, with solid and dashed lines indicating the mean performance and shaded areas representing one standard deviation.

Except for the correction strength L in SoCo, all hyperparameters are identical to the default settings. For evaluation, each trained policy is tested over 40 episodes, and the average return is reported. All experiments are repeated with 3 random seeds to account for variance. Detailed hyperparameter settings are provided in Appendix D.4.

#### 4.2 EVALUATION RESULTS AND ANALYSIS

We evaluate SoCo on nine tasks across four scenarios with varying characteristics and difficulty. Across both backbone algorithms, SoCo improves training efficiency and achieves competitive or superior performance, demonstrating its effectiveness.

**Spread.** In the *Spread* tasks, agents must learn not only to navigate to landmarks but also to resolve target assignment and avoid collisions, with difficulty growing rapidly as the number of agents increases. As shown in Figures 2a–2c, training from scratch becomes highly inefficient under this setting. With SoCo, however, agents first acquire basic navigation skills from solo demonstrations, and during cooperative training, they only need to master target selection and collision avoidance via policy fusion. This significantly improves both training efficiency and final performance. For example, in the 5-agent task, SoCo converges faster and outperforms both backbone algorithms by more than 20% in final performance, demonstrating its effectiveness in mitigating the challenge of goal ambiguity.

**LongSwimmer.** In the *LongSwimmer* task, agents collaboratively control a multi-segment worm to swim forward. As shown in Figures 2d-2f, although the control difficulty is moderate and both backbone algorithms and SoCo eventually reach similar performance, our framework effectively speeds up training. For example, in the 3-agent task, HATD3-SoCo attains an average return of about 300 at roughly 1.0M steps, while vanilla HATD3 only reaches a comparable level around

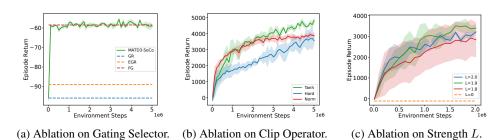


Figure 3: Ablation study of SoCo. Results are averaged over three random seeds, with solid and dashed lines indicating the mean performance and shaded areas representing one standard deviation.

1.6M steps, saving nearly **40%** of training samples. These results highlight that SoCo successfully leverages solo demonstrations as a scalable prior to accelerate cooperative learning.

**MultiHalfcheetah.** The *MultiHalfCheetah* task requires multiple HalfCheetah agents connected by elastic tendons to run forward in coordination. The tendon coupling already introduces instability, while the intrinsic difficulty of HalfCheetah control further compounds the challenge. Unlike *Spread* or *LongSwimmer*, this scenario also alters the agents' mass, creating domain shifts that make solo policies non-transferable. Nevertheless, SoCo leverages action editor to adapt solo priors to the shifted dynamics while retaining their basic control skills. As shown in Figures 2g and 2h, this leads to markedly improved training efficiency for the backbone algorithms. In particular, on the 3-agent task, HATD3-SoCo improves the final performance by approximately **83.91**% over the backbone, highlighting the strength of SoCo in leveraging solo demonstrations to boost both efficiency and effectiveness of cooperative training.

**MultiWalker.** The MultiWalker task is the most challenging among the four scenarios. Agents must not only stabilize multiple walkers but also coordinate to carry a long, unstable package under noisy observations. The reward structure is harsh, and the domain shifts are severe, making direct transfer highly difficult. In this setting, backbone MARL algorithms struggle to learn effective package transport within the training budget. By contrast, SoCo leverages policy fusion to refine solo priors and adapt them to unstable multi-agent dynamics, enabling faster discovery of transport strategies and yielding clear gains in both training speed and final performance. Notably, as shown in 2i, SoCo improves the final performance by **91.51%** on HATD3 and **11.97%** on MATD3 compared to their vanilla counterparts. This shows that SoCo can transfer solo knowledge even under extreme conditions, substantially improving both training efficiency and performance.

#### 4.3 ABLATION STUDY

#### 4.3.1 COMPONENT ABLATION

Gating Selector. We conduct an ablation study on the 3-agent Spread task to isolate the effect of the gating selector. Given the environment structure, setting the correction strength to zero (L=0) already yields strong performance, so we focus exclusively on the gating component. We compare three variants: (i) **Random Gating (RG)**, where targets are sampled randomly at each step; (ii) **Episode-wise Random Gating (ERG)**, where targets are randomly fixed at the start of each episode; and (iii) **Fixed Gating (FG)**, where *distinct* targets are deterministically assigned by agent index at the beginning of each episode, serving as an oracle assignment in this scenario. As shown in Figure 3a, randomized gating (RG / ERG) suffers from frequent target conflicts and poor coordination, whereas our learned gating selector can avoid conflicts and perform competitively to FG.

**Clip Operator.** As discussed in Section 3.3.2, we adopt a tanh-based clip operator to prevent fused actions from exceeding valid ranges. Nevertheless, SoCo is designed as a general framework, and different clipping strategies can be customized for specific tasks. To examine this flexibility and assess the suitability of our choice, we conduct experiments on the 2-agent *MultiHalfCheetah* task, evaluating how alternative operators affect both training efficiency and final performance. Using MATD3 as the backbone, we compare two variants:

- (i) **Norm**, which normalizes the action as  $\text{Clip}(\tilde{a}_t^i + \Delta a_t^i) = (\tilde{a}_t^i + \Delta a_t^i)/(L+1)$ ;
- (ii) **Hard**, which directly truncates actions via clamp( $\tilde{a}_t^i + \Delta a_t^i, -1, 1$ ).

As shown in Figure 3b, the Norm operator accelerates early learning but suffers from weak asymptotic performance, as normalization continuously shrinks the effective action magnitude and reduces policy expressiveness. The Hard operator, on the other hand, truncates actions abruptly, suppressing gradient signals and leading to slow and unstable training. In contrast, our tanh-based design achieves a smoother balance between boundedness and gradient flow, since gradients are only compressed near the action boundaries. This enables both stable learning dynamics and stronger final performance, making the tanh-based operator a natural and effective default choice for SoCo. That said, the framework remains flexible to alternative operators when required by task dynamics.

# 4.3.2 Hyperparameter Sensitivity

An important hyperparameter in the SoCo framework is the correction strength L, which controls the degree to which the algorithm leverages knowledge from solo demonstrations. We conduct experiments on the 2-agent MultiHalfCheetah task with  $L \in \{0, 1.8, 1.9, 2.0\}$ . Since this environment does not involve multi-goal settings, we can effectively isolate the influence of the gating selector and focus on the impact of this hyperparameter on SoCo's performance. The results in Table 3c show that when L is too small, SoCo relies excessively on solo demonstrations, which limits its training efficiency. In contrast, when L is too large, SoCo adapts quickly to environmental changes in the early stage, but insufficient use of solo knowledge makes it difficult to discover better strategies later, leading to suppressed final performance. Thus, an appropriate choice of L is essential for maximizing the effectiveness of SoCo. In addition, we also examined the case of L=0, where solo policies are directly applied in the multi-agent environment. The results reveal that domain shift prevents the agents from being successfully controlled, underscoring the necessity of policy fusion.

# 5 RELATED WORK

MARL. Multi-Agent Reinforcement Learning (MARL) has advanced rapidly, giving rise to diverse paradigms. Fully decentralized methods train and execute policies without centralized information (Tampuu et al., 2017; de Witt et al., 2020), but often suffer from limited coordination. By contrast, the Centralized Training with Decentralized Execution (CTDE) paradigm (Ackermann et al., 2019; Rashid et al., 2020; Yu et al., 2022; Zhong et al., 2024; Li et al., 2025b) has become dominant, enabling centralized training for coordination while preserving decentralized execution. In this paper, we focus on deterministic policy gradient methods under CTDE.

**Transferable MARL.** To mitigate the high cost of training from scratch, transferable MARL aims to reuse experience from source tasks to accelerate learning in target tasks with limited interaction. Existing approaches include offline-to-online (Zhong et al., 2025), multi-task (Chen et al., 2024; Liu et al., 2025; Jha et al., 2025), ad-hoc teamwork (Zhang et al., 2023b; Li et al., 2025a), and mixed-component or personalized data (Wang et al., 2023; Yu et al., 2025). While these methods broaden MARL's applicability, they still assume sufficient multi-agent data. In contrast, exploiting solo demonstrations, abundant yet lacking cooperative signals, remains underexplored. Our work fills this gap by showing that such data can be effectively leveraged to improve both the training efficiency and performance of cooperative learning.

More detailed discussions are provided in Appendix B.

# 6 Conclusion

In this paper, we studied an underexplored problem: how to exploit solo demonstrations to accelerate MARL. We propose a novel Solo-to-Collaborative RL (SoCo) framework, which leverages solo demonstrations by pretraining a shared solo policy and adapting it during cooperative training through policy fusion with a gating selector and an action editor. Experiments across diverse tasks show that SoCo improves training efficiency and achieves competitive or even superior performance, highlighting that solo demonstrations provide a scalable and effective complement to multi-agent data, making cooperative learning more practical and broadly applicable.

# REPRODUCIBILITY STATEMENT

To ensure reproducibility, we detail our experimental setup in Section 4 and Appendix D, covering environment configuration, solo demonstration collection, implementation details, and hyperparameters. The source code will be made publicly available upon publication.

# REFERENCES

- Johannes Ackermann, Volker Gabler, Takayuki Osa, and Masashi Sugiyama. Reducing overestimation bias in multi-agent domains using double centralized critics, 2019. URL https://arxiv.org/abs/1910.01465.
- Christopher Amato. An Introduction to Centralized Training for Decentralized Execution in Cooperative Multi-Agent Reinforcement Learning, 2024. URL https://arxiv.org/abs/2409.03052.
- Lucian Busoniu, Robert Babuska, and Bart De Schutter. A comprehensive survey of multiagent reinforcement learning. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 38(2):156–172, 2008.
- Weilin Cai, Juyong Jiang, Fan Wang, Jing Tang, Sunghun Kim, and Jiayi Huang. A survey on mixture of experts in large language models. *IEEE Transactions on Knowledge and Data Engineering*, 2025.
- Xiao Cao, Mingyang Li, Yuting Tao, and Peng Lu. Hma-sar: Multi-agent search and rescue for unknown located dynamic targets in completely unknown environments. *IEEE Robotics and Automation Letters*, 9(6):5567–5574, 2024. doi: 10.1109/LRA.2024.3396097.
- Jiayu Chen, Tian Lan, and Vaneet Aggarwal. Variational offline multi-agent skill discovery. *arXiv* preprint arXiv:2405.16386, 2024.
- Rodrigo de Lazcano, Kallinteris Andreas, Jun Jet Tai, Seungjae Ryan Lee, and Jordan Terry. Gymnasium robotics, 2024. URL http://github.com/Farama-Foundation/Gymnasium-Robotics.
- Christian Schroeder de Witt, Tarun Gupta, Denys Makoviichuk, Viktor Makoviychuk, Philip HS Torr, Mingfei Sun, and Shimon Whiteson. Is independent learning all you need in the starcraft multi-agent challenge? *arXiv preprint arXiv:2011.09533*, 2020.
- Perry Dong, Qiyang Li, Dorsa Sadigh, and Chelsea Finn. Expo: Stable reinforcement learning with expressive policies. *arXiv preprint arXiv:2507.07986*, 2025a.
- Yihong Dong, Xue Jiang, Jiaru Qian, Tian Wang, Kechi Zhang, Zhi Jin, and Ge Li. A survey on code generation with llm-based agents, 2025b. URL https://arxiv.org/abs/2508.00083.
- Scott Fujimoto, Herke van Hoof, and David Meger. Addressing function approximation error in actor-critic methods. In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1587–1596. PMLR, 10–15 Jul 2018. URL https://proceedings.mlr.press/v80/fujimoto18a.html.
- Sven Gronauer and Klaus Diepold. Multi-agent deep reinforcement learning: a survey. *Artificial Intelligence Review*, 55(2):895–943, 2022.
- Jayesh K Gupta, Maxim Egorov, and Mykel Kochenderfer. Cooperative multi-agent control using deep reinforcement learning. In *International Conference on Autonomous Agents and Multiagent Systems*, pp. 66–83. Springer, 2017.
  - Pablo Hernandez-Leal, Bilal Kartal, and Matthew E Taylor. A survey and critique of multiagent deep reinforcement learning. *Autonomous Agents and Multi-Agent Systems*, 33(6):750–797, 2019.

- Siyi Hu, Fengda Zhu, Xiaojun Chang, and Xiaodan Liang. UPDeT: Universal Multi-agent RL via Policy Decoupling with Transformers. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=v9c7hr9ADKx.
  - Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *International Conference on Learning Representations*, 2017. URL https://openreview.net/forum?id=rkE3y85ee.
  - Kunal Jha, Wilka Carvalho, Yancheng Liang, Simon Shaolei Du, Max Kleiman-Weiner, and Natasha Jaques. Cross-environment cooperation enables zero-shot multi-agent coordination. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=zBBYsVGKuB.
  - Shyam Sundar Kannan, Vishnunandan L. N. Venkatesh, and Byung-Cheol Min. Smart-Ilm: Smart multi-agent robot task planning using large language models. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 12140–12147, 2024. doi: 10.1109/IROS58592.2024.10802322.
  - Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
  - Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
  - Andreas Kontogiannis, Konstantinos Papathanasiou, Yi Shen, Giorgos Stamou, Michael M. Zavlanos, and George Vouros. Enhancing cooperative multi-agent reinforcement learning with state modelling and adversarial exploration. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=TCsdlqzZNL.
  - Lihe Li, Lei Yuan, Pengsen Liu, Tao Jiang, and Yang Yu. LLM-assisted semantically diverse teammate generation for efficient multi-agent coordination. In *Forty-second International Conference on Machine Learning*, 2025a. URL https://openreview.net/forum?id=Vhktpw6Vid.
  - Yueheng Li, Guangming Xie, and Zongqing Lu. Revisiting cooperative off-policy multi-agent reinforcement learning. In *Forty-second International Conference on Machine Learning*, 2025b. URL https://openreview.net/forum?id=JPkJAyutW0.
  - Zhuoran Li, Ling Pan, and Longbo Huang. Beyond conservatism: Diffusion policies in offline multi-agent reinforcement learning. *arXiv preprint arXiv:2307.01472*, 2023.
  - Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv* preprint arXiv:2210.02747, 2022.
  - Iou-Jen Liu, Raymond A. Yeh, and Alexander G. Schwing. Pic: Permutation invariant critic for multi-agent deep reinforcement learning. In Leslie Pack Kaelbling, Danica Kragic, and Komei Sugiura (eds.), Proceedings of the Conference on Robot Learning, volume 100 of Proceedings of Machine Learning Research, pp. 590–602. PMLR, 30 Oct-01 Nov 2020. URL https:// proceedings.mlr.press/v100/liu20a.html.
  - Jiarong Liu, Yifan Zhong, Siyi Hu, Haobo Fu, QIANG FU, Xiaojun Chang, and Yaodong Yang. Maximum Entropy Heterogeneous-Agent Reinforcement Learning. In *The Twelfth International Conference on Learning Representations*, 2024a. URL https://openreview.net/forum?id=tmqOhBC4a5.
  - Sicong Liu, Yang Shu, Chenjuan Guo, and Bin Yang. Learning Generalizable Skills from Offline Multi-Task Data for Multi-Agent Cooperation. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=HRlujVR0ig.
    - Zongkai Liu, Qian Lin, Chao Yu, Xiawei Wu, Yile Liang, Donghui Li, and Xuetao Ding. Offline Multi-Agent Reinforcement Learning via In-Sample Sequential Policy Optimization, 2024b. URL https://arxiv.org/abs/2412.07639.

- Ryan Lowe, YI WU, Aviv Tamar, Jean Harb, OpenAI Pieter Abbeel, and Igor Mordatch. Multiagent actor-critic for mixed cooperative-competitive environments. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper\_files/paper/2017/file/68a9750337a418a86fe06c1991ald64c-Paper.pdf.
- Laetitia Matignon, Laurent Jeanpierre, and Abdel-Illah Mouaddib. Coordinated Multi-Robot Exploration Under Communication Constraints Using Decentralized Markov Decision Processes. *Proceedings of the AAAI Conference on Artificial Intelligence*, 26(1):2017–2023, Sep. 2021. doi: 10. 1609/aaai.v26i1.8380. URL https://ojs.aaai.org/index.php/AAAI/article/view/8380.
- Frans A Oliehoek, Matthijs TJ Spaan, and Nikos Vlassis. Optimal and approximate Q-value functions for decentralized POMDPs. *Journal of Artificial Intelligence Research*, 32:289–353, 2008.
- Frans A Oliehoek, Christopher Amato, et al. *A concise introduction to decentralized POMDPs*, volume 1. Springer, 2016.
- Ling Pan, Longbo Huang, Tengyu Ma, and Huazhe Xu. Plan Better Amid Conservatism: Offline Multi-Agent Reinforcement Learning with Actor Rectification. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 17221–17237. PMLR, 17–23 Jul 2022. URL https://proceedings.mlr.press/v162/pan22a.html.
- Bei Peng, Tabish Rashid, Christian Schroeder de Witt, Pierre-Alexandre Kamienny, Philip Torr, Wendelin Boehmer, and Shimon Whiteson. Facmac: Factored multi-agent centralised policy gradients. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan (eds.), Advances in Neural Information Processing Systems, volume 34, pp. 12208–12221. Curran Associates, Inc., 2021. URL https://proceedings.neurips.cc/paper\_files/paper/2021/file/65b9eea6e1cc6bb9f0cd2a47751a186f-Paper.pdf.
- Tabish Rashid, Mikayel Samvelyan, Christian Schroeder de Witt, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. Monotonic Value Function Factorisation for Deep Multi-Agent Reinforcement Learning. *Journal of Machine Learning Research*, 21(178):1–51, 2020. URL http://jmlr.org/papers/v21/20-081.html.
- Jianzhun Shao, Yun Qu, Chen Chen, Hongchang Zhang, and Xiangyang Ji. Counterfactual Conservative Q Learning for Offline Multi-agent Reinforcement Learning. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=62zmO4mv8X.
- Yoav Shoham and Kevin Leyton-Brown. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations*. Cambridge University Press, 2008.
- EP Stepanov, RL Smeliansky, AV Plakunov, AV Borisov, Xia Zhu, Jianing Pei, and Zhen Yao. On fair traffic allocation and efficient utilization of network resources based on marl. *Computer Networks*, 250:110540, 2024.
- Peter Stone, Gal Kaminka, Sarit Kraus, and Jeffrey Rosenschein. Ad hoc autonomous agent teams: Collaboration without pre-coordination. *Proceedings of the AAAI Conference on Artificial Intelligence*, 24(1):1504–1509, Jul. 2010. doi: 10.1609/aaai.v24i1.7529. URL https://ojs.aaai.org/index.php/AAAI/article/view/7529.
- Shaoqi Sun, Yuanzhao Zhai, Kele Xu, Dawei Feng, and Bo Ding. Progressive diversifying policy for multi-agent reinforcement learning. In *ICASSP 2023 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5, 2023. doi: 10.1109/ICASSP49357. 2023.10096125.
- Ardi Tampuu, Tambet Matiisen, Dorian Kodelja, Ilya Kuzovkin, Kristjan Korjus, Juhan Aru, Jaan Aru, and Raul Vicente. Multiagent cooperation and competition with deep reinforcement learning. *PLOS ONE*, 12(4):1–15, 04 2017. doi: 10.1371/journal.pone.0172395. URL https://doi.org/10.1371/journal.pone.0172395.

- Chen Tang, Ben Abbatematteo, Jiaheng Hu, Rohan Chandra, Roberto Martín-Martín, and Peter Stone. Deep reinforcement learning for robotics: A survey of real-world successes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 28694–28698, 2025.
  - J Terry, Benjamin Black, Nathaniel Grammel, Mario Jayakumar, Ananth Hari, Ryan Sullivan, Luis S Santos, Clemens Dieffendahl, Caroline Horsch, Rodrigo Perez-Vicente, et al. Pettingzoo: Gym for multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 34: 15032–15043, 2021.
  - Caroline Wang, Ishan Durugkar, Elad Liebman, and Peter Stone. Dm<sup>2</sup>: Decentralized multi-agent reinforcement learning via distribution matching. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 11699–11707, 2023.
  - Tonghan Wang, Tarun Gupta, Anuj Mahajan, Bei Peng, Shimon Whiteson, and Chongjie Zhang. RODE: Learning Roles to Decompose Multi-Agent Tasks. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=TTUVg6vkNjK.
  - Jizhou Wu, Jianye Hao, Tianpei Yang, Xiaotian Hao, Yan Zheng, Weixun Wang, and Matthew E. Taylor. PORTAL: Automatic Curricula Generation for Multiagent Reinforcement Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(14):15934–15942, Mar. 2024. doi: 10.1609/aaai.v38i14.29524. URL https://ojs.aaai.org/index.php/AAAI/article/view/29524.
  - Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM computing surveys*, 56(4):1–39, 2023.
  - Chao Yu, Akash Velu, Eugene Vinitsky, Jiaxuan Gao, Yu Wang, Alexandre Bayen, and YI WU. The Surprising Effectiveness of PPO in Cooperative Multi-Agent Games. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), Advances in Neural Information Processing Systems, volume 35, pp. 24611–24624. Curran Associates, Inc., 2022. URL https://proceedings.neurips.cc/paper\_files/paper/2022/file/9c1535a02f0ce079433344e14d910597-Paper-Datasets\_and\_Benchmarks.pdf.
  - Peihong Yu, Manav Mishra, Alec Koppel, Carl Busart, Priya Narayan, Dinesh Manocha, Amrit Singh Bedi, and Pratap Tokekar. Beyond joint demonstrations: Personalized expert guidance for efficient multi-agent reinforcement learning. *Transactions on Machine Learning Research*, 2025. ISSN 2835-8856. URL https://openreview.net/forum?id=kzPNHQ8ByY.
  - Xianghua Zeng, Hang Su, Zhengyi Wang, and Zhiyuan LIN. Graph diffusion for robust multiagent coordination. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=T5IZ32ImAB.
  - Fuxiang Zhang, Chengxing Jia, Yi-Chen Li, Lei Yuan, Yang Yu, and Zongzhang Zhang. Discovering Generalizable Multi-agent Coordination Skills from Multi-task Offline Data. In *The Eleventh International Conference on Learning Representations*, 2023a. URL https://openreview.net/forum?id=53FyUAdP7d.
  - Ruiqi Zhang, Jing Hou, Florian Walter, Shangding Gu, Jiayi Guan, Florian Röhrbein, Yali Du, Panpan Cai, Guang Chen, and Alois Knoll. Multi-agent reinforcement learning for autonomous driving: A survey. *arXiv* preprint arXiv:2408.09675, 2024.
  - Ziqian Zhang, Lei Yuan, Lihe Li, Ke Xue, Chengxing Jia, Cong Guan, Chao Qian, and Yang Yu. Fast Teammate Adaptation in the Presence of Sudden Policy Change. In Robin J. Evans and Ilya Shpitser (eds.), *Proceedings of the Thirty-Ninth Conference on Uncertainty in Artificial Intelligence*, volume 216 of *Proceedings of Machine Learning Research*, pp. 2465–2476. PMLR, 31 Jul–04 Aug 2023b. URL https://proceedings.mlr.press/v216/zhang23a.html.
  - Hai Zhong, Xun Wang, Zhuoran Li, and Longbo Huang. Offline-to-online multi-agent reinforcement learning with offline value function memory and sequential exploration. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems*, pp. 2373–2381, 2025.

# **Technical Appendices**

A LLM Usage **Detailed Related Work** Algorithm Pseudocode **D** Experiments Details 

# A LLM USAGE

In this paper, we use the LLM to polish writing and check grammar issues.

# B DETAILED RELATED WORK

MARL. Multi-Agent Reinforcement Learning (MARL) has advanced rapidly in recent years, giving rise to diverse paradigms and methods. Fully decentralized approaches train and execute policies without centralized information (Tampuu et al., 2017; de Witt et al., 2020), but their performance is often constrained by the absence of communication among agents. By contrast, the Centralized Training with Decentralized Execution (CTDE) paradigm (Oliehoek et al., 2008; Matignon et al., 2021; Amato, 2024; Li et al., 2025b) has become the mainstream, enabling agents to learn with centralized information for coordination while still executing policies in a decentralized manner. Representative algorithms include HASAC (Liu et al., 2024a), HARL (Zhong et al., 2024), MAPPO (Yu et al., 2022), QMIX (Rashid et al., 2020), and MATD3 (Ackermann et al., 2019). In this work, we adopt deterministic policy gradient methods within the CTDE paradigm, with particular focus on MATD3 and HATD3.

**Transferable MARL.** Since training MARL from scratch is often sample-inefficient and costly, transferable MARL seeks to reuse experience from source settings to accelerate learning in target tasks with limited additional interaction. Existing approaches span several directions: offline-toonline MARL (Zhong et al., 2025), which leverages offline pretraining to speed up online exploration and correct distributional shift; multi-task MARL (Hu et al., 2021; Wang et al., 2021; Zhang et al., 2023a; Chen et al., 2024; Liu et al., 2025; Jha et al., 2025), which extracts transferable knowledge from multiple source tasks and applies it to unseen ones; ad-hoc teamwork (Stone et al., 2010; Zhang et al., 2023b; Li et al., 2025a), which exposes agents to diverse teammates to improve robustness when coordinating with unseen partners; and MARL with mixed-component or personalized data (Wang et al., 2023; Yu et al., 2025), which constructs datasets from individual trajectories generated by different cooperative policies, enriching training diversity while preserving per-step consistency. While these methods broaden the applicability of MARL, they all rely on sufficient multi-agent data. In contrast, the potential of exploiting solo demonstrations, abundant but lacking cooperative signals, remains largely unexplored. Our work addresses this gap by showing that such data can be effectively leveraged to accelerate cooperative training, thereby opening a promising new avenue.

# C ALGORITHM PSEUDOCODE

# Algorithm 1 Solo-to-Collaborative Reinforcement Learning (SoCo) **Input:** Datasets of solo demonstration $\mathcal{D}$ and edit strength L. Initialize the parameters w for solo policy $\beta_w$ , $\phi = \{\varphi, \theta\}$ for weight assigner $g_{\varphi}$ and coordination policy $\pi_{\theta}$ , $\{\psi_j\}_{j=1}^2$ for $\{Q_j\}_{j=1}^2$ , and $\bar{\psi}_1, \bar{\psi}_2, \bar{\phi}$ for target networks. Train solo policy $\beta_w$ with $\mathcal{D}$ according to Eq. (1). Initialize the replay buffer $\mathcal{B}$ . for i=1 to $T_{\rm max}$ do Obtain the joint observation $\mathbf{o}_t$ from the environment. // Agent-wise Solo-to-Collaborative Transfer for n=1 to N do // Observation Decomposition Decompose local observation $o_t^n$ into solo views $\{o_t^{n,k}\}_{k=1}^{G_n}$ . // Policy Fusion Calculate $a_t^n = \beta_w(\{o_t^n\})$ and obtain solo action $\tilde{a}_t^n$ by Eq. (3) Calculate editing action $\Delta a_t^n$ by Eq. (4). Obtain final action $a_t^n$ by combining $\tilde{a}_t^n$ and $\Delta a_t^n$ according to Eq. (5). end for // Cooperative MARL Training Use $a_t = (a_t^1, \dots, a_t^N)$ to interact with the environment and save $(s_t, o_t, a_t, r_t, o_{t+1})$ into $\mathcal{B}$ . Sample a batch of transitions $\{(s_t, o_t, a_t, r_t, o_{t+1})\}$ from $\mathcal{B}$ .

Update critics  $Q_1$ ,  $Q_2$  and fused policy  $\Pi_{\phi}$  through standard MARL algorithms.

#### D EXPERIMENTS DETAILS

#### D.1 Environments

end for

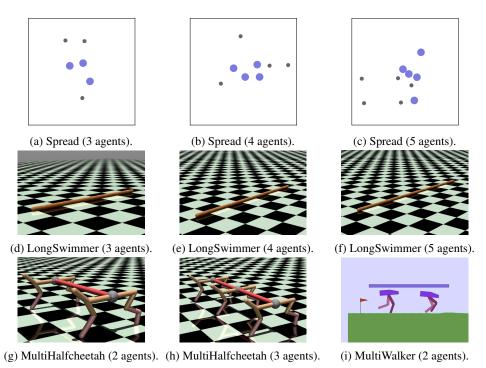


Figure 4: All the cooperative tasks in our experiments.

We evaluate SoCo on nine tasks across four representative cooperative scenarios:

 **Spread** (Lowe et al., 2017; Terry et al., 2021). As shown in Figures 4a-4c, in this environment, N agents are initialized at random positions in a bounded 2D plane, while K=N landmarks are also randomly placed without overlap. Agents must navigate to distinct landmarks while avoiding collisions. The per-step reward for each agent i is defined as the average of a global and a local component:

$$r_t^i = \frac{1}{2} \left( r_t^{\text{global}} + r_t^{\text{local},i} \right).$$

The global reward is shared across agents and encourages coverage of landmarks:

$$r_t^{\text{global}} = -\sum_{k=1}^K \min_{j \in \mathcal{N}} \|p_t^j - l_k\|_2,$$

where  $p_t^j$  is the position of agent j, and  $l_k$  is the position of landmark k.

The local reward penalizes collisions:

$$r_t^{\text{local},i} = \begin{cases} -C_t^i, & \text{if agent } i \text{ collides with } C_t^i \text{ other agents,} \\ 0, & \text{otherwise.} \end{cases}$$

Finally, the environment reward is the sum over all agents' individual rewards:

$$R_t = \sum_{i \in \mathcal{N}} r_t^i.$$

We evaluate on tasks with 3, 4, and 5 agents.

**LongSwimmer** (Peng et al., 2021; de Lazcano et al., 2024). As shown in Figures 4d and 4f, in this environment, a (2N+1)-segment worm must be controlled to swim forward. Each pair of adjacent segments is connected by a joint, and each agent is responsible for controlling two consecutive joints in sequence. The worm's initial state is sampled from a uniform distribution within a predefined range, while its initial velocity is drawn from Gaussian noise to diversify the dynamics. The per-step reward for each agent i is:

$$r_t^i = v_t - 0.0001 \cdot \sum_{i \in \mathcal{N}} ||a_t^i||_2^2,$$

where  $v_t$  is the forward velocity of the worm,  $a_t^i$  is the action taken by agent i.

The environment reward is defined as the average of all agents' rewards:

$$R_t = \frac{1}{N} \sum_{i \in \mathcal{N}} r_t^i.$$

We evaluate on tasks with 3, 4, and 5 agents.

MultiHalfCheetah (Peng et al., 2021; de Lazcano et al., 2024). As shown in Figures 4g and 4h, in this environment, N HalfCheetah agents are connected in series by elastic tendons and must collaboratively run forward. Each agent's initial state is sampled from a uniform distribution within a predefined range, and its initial velocity is drawn from Gaussian noise to diversify dynamics. The per-step reward for each agent i is:

$$r_t^i = v_t^i - 0.1 \cdot ||a_t^i||_2^2,$$

where  $v_t^i$  is the forward velocity of agent i,  $a_t^i$  is its action.

The environment reward is defined as the average of all agents' rewards:

$$R_t = \frac{1}{N} \sum_{i \in \mathcal{N}} r_t^i.$$

We evaluate on tasks with 2 and 3 HalfCheetahs.

MultiWalker (Gupta et al., 2017; Terry et al., 2021). As shown in Figure 4i, in this environment, N bipedal robots must collaboratively lift and carry a long package forward. The terrain has a randomly undulating profile at the start of each episode. Walkers are initialized at fixed, equally spaced positions in standing poses; to diversify initial conditions, a small random external force is applied to each walker's head at t=0. The package length scales proportionally with the number of walkers, and each walker's observation is corrupted with noise.

At each step, each walker receives a progress reward equal to the forward displacement of the package, plus a small shaping penalty for head tilting and a -10 penalty if a walker falls:

$$r_t^i \ = \ \Delta x_t^{\mathrm{package}} \ - \ 5 \cdot \Delta \theta_t^{\mathrm{head},i} \ - \ 10 \cdot \mathbf{1} \{ \mathrm{walker} \ i \ \mathrm{falls} \}.$$

Episodes terminate if the package falls, leaves the left edge, or if any walker falls, in which case all walkers receive -100. If the package exits the right edge, termination occurs with reward 0.

The environment reward at each step is the sum of individual rewards:

$$R_t = \sum_{i \in \mathcal{N}} r_t^i.$$

We evaluate on task with 2 walkers.

#### D.2 SOLO DEMONSTRATION

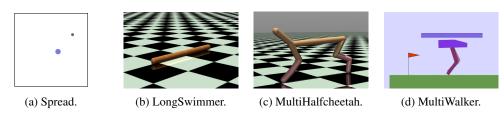


Figure 5: Solo tasks corresponding to each cooperative scenario.

# D.2.1 DATA COLLECTION

For each cooperative scenario, we first train a policy on its corresponding solo task using TD3 (Fujimoto et al., 2018), and then collect 1M transitions to learn the solo policy. Table 1 summarizes the average episode returns of the solo task demonstrations.

Table 1: Average episode return of collected solo demonstrations.

Scenario	Spread	LongSwimmer	MultiHalfcheetah	MultiWalker
Average Episode Return	-14.77	119.44	7054.87	197.18

#### D.2.2 SOLO TASKS VS. COOPERATIVE SCENARIOS

These solo tasks, illustrated in Figure 5, exhibit noticeable gaps from their cooperative counterparts, ranging from goal ambiguity (*Spread*), to moderate domain shift (*LongSwimmer*), to notable domain

shift and cooperative difficulty (*MultiHalfCheetah*), and to severe domain shift with substantial cooperative difficulty (*MultiWalker*).

Specifically, in the *Spread* scenario, the solo task allows an agent to observe only a single target, whereas in the cooperative setting, multiple targets are visible simultaneously. In the *LongSwimmer*, the motion of the worm is affected by the actions of other agents, introducing a moderate domain shift. In *MultiHalfCheetah*, the solo task doubles the agent's mass and removes tendon constraints, making it simpler than the coupled cooperative case. Finally, in *MultiWalker*, the solo task differs drastically from the cooperative environment: the package length and walker positions change, observations are noisy, and interference from teammates is absent in solo but present in multi-agent training, resulting in severe domain shift and substantially higher cooperative difficulty.

#### D.3 IMPLEMENTATION DETAILS

 Our implementation and experiments are based on the HARL codebase (Zhong et al., 2024). The additional components introduced by SoCo, i.e., the solo policy, gating selector, and action editor, share the same architecture as the backbone actor network, implemented as 2-layer MLPs with ReLU activations. For action fusion, we adopt a tanh-based clip operator: when  $\Delta a \equiv 0$ , no constraint is applied; otherwise, the fused action is bounded through a tanh transformation. We use Adam (Kingma, 2014) for optimization. Additionally, in the 3-agent *MultiHalfCheetah* environment, the tendon structure can destabilize the MuJoCo simulator. To mitigate this, we impose an additional constraint on the output of HATD3-SoCo, clipping it to the range [-0.85, 0.85].

#### D.4 HYPERPARAMETERS

Except for the correction strength L in SoCo, all hyperparameters follow the default or recommended (when available) settings in HARL to ensure fair comparison. The detailed configurations are reported in Table 2.

Table 2: Shared hyperparameters for all algorithms.

Hyperparameter	Value	Hyperparameter	Value
Batch Size	1000	Buffer Size	1000000
Hidden Size	256 (128 for <i>Spread</i> )	Discount Factor $\gamma$	0.99
n-step TD	10 (1 for <i>Spread</i> )	Explore Noise	0.1
Policy Noise	0.2	Noise Clip	0.5
Policy Delay	2	Soft Update Coefficient	0.005
Actor Learning Rate	0.0005	Critic Learning Rate	0.001

For SoCo, L is an important hyperparameter that controls the extent to which knowledge from solo demonstrations is leveraged. The values of L used for each task and backbone algorithm are summarized in Table 3. Different tasks require different L values, as the optimal balance depends on factors such as the degree of domain shift and the inherent difficulty of the cooperative environment.

Table 3: Correction strength L used in SoCo for each task and backbone algorithm.

Task	MATD3-SoCo	HATD3-SoCo
Spread-3	0	0
Spread-4	0	0
Spread-5	0	0
LongSwimmer-3	3.15	2.20
LongSwimmer-4	3.10	2.90
LongSwimmer-5	2.10	2.85
MultiHalfCheetah-2	1.90	2.00
MultiHalfCheetah-3	1.90	1.90
MultiWalker-2	2.00	2.00