

PADiff: Predictive and Adaptive Diffusion Policies for Ad Hoc Teamwork

Hohei Chan¹, Xinzhi Zhang¹, Antao Xiang¹, Weinan Zhang², Mengchen Zhao^{*1}

¹School of Software Engineering, South China University of Technology, Guangzhou, China

²Shanghai Jiao Tong University, Shanghai, China

hoheichanchx@gmail.com, zhang.xinzhi@outlook.com, 202330421721@mail.scut.edu.cn, wnzhang@sjtu.edu.cn, zzmc@scut.edu.cn

Abstract

Ad hoc teamwork (AHT) requires agents to collaborate with previously unseen teammates, which is crucial for many real-world applications. The core challenge of AHT is to develop an ego agent that can predict and adapt to unknown teammates on the fly. Conventional RL-based approaches optimize a single expected return, which often causes policies to collapse into a single dominant behavior, thus failing to capture the multimodal cooperation patterns inherent in AHT. In this work, we introduce PADiff, a diffusion-based approach that captures agent’s multimodal behaviors, unlocking its diverse cooperation modes with teammates. However, standard diffusion models lack the ability to predict and adapt in highly non-stationary AHT scenarios. To address this limitation, we propose a novel diffusion-based policy that integrates critical predictive information about teammates into the denoising process. Extensive experiments across three cooperation environments demonstrate that PADiff outperforms existing AHT methods significantly.

Introduction

Ad hoc teamwork (AHT) presents a core challenge in multi-agent systems (Stone et al. 2010), which requires agents to collaborate with unknown teammates without predefined coordination (Yuan et al. 2023a). Consider a robotic soccer match, where an autonomous agent plays with unfamiliar teammates. The agent must make real-time decisions without any knowledge of the new teammates’ playing styles or strategies. These challenges are not limited to simulated environments. In real-world applications like disaster response (Yucesoy, Balcik, and Coban 2025), agents must coordinate with human partners operating under unknown procedures. In autonomous driving (Teng et al. 2023), vehicles often need to interact with other road users without established communication protocols. Such situations require agents to infer teammate behaviors, adapt their strategies, and cooperate effectively on the fly (Ravula 2019). As multi-agent systems are deployed in real world (Yuan et al. 2023b), it is urgent to address the challenges of AHT and improve the autonomy of multi-agent systems.

AHT fundamentally requires an ego agent to collaborate with previously unseen teammates, making it crucial to prepare for multiple possible cooperative patterns rather than relying on a single optimal response. Consider a robotic soccer scenario involving AHT shown in Figure 1): When an ego agent advances with the ball, she has to choose between two modes of cooperation: either assist a teammate or directly attempt a shot. This example indicates that the ego agent must learn multiple cooperation patterns simultaneously, highlighting the multimodal nature of the optimal policies. However, most existing solutions to AHT are based on reinforcement learning (RL) (Chen et al. 2020; Gu et al. 2021; Papoudakis, Christianos, and Albrecht 2021), which often fail to learn multimodal policies. RL typically optimizes for a fixed expected return, causing the policy to collapse to a single dominant behavior. From the generalization perspective, even if RL learns an optimal policy, it often still fails to generalize when the teammates change their cooperation patterns. Although previous AHT methods (e.g., LIAM(Papoudakis, Christianos, and Albrecht 2021) and GPL-SPI(Rahman et al. 2021)) leverage maximum entropy RL (e.g., SAC) to enhance stochastic exploration and mitigate policy collapse into a singular cooperative mode, their reliance on undirected behavioral dispersion inherently fails to structurally model multimodal distributions, thus poorly capturing diverse cooperation strategies. In contrast, diffusion models (Croitoru et al. 2023) are naturally suitable to model multimodal distributions(Wang, Hunt, and Zhou 2022), making them stronger potential solutions to AHT, where learning multiple cooperation patterns simultaneously is essential for robust teamwork.

Diffusion models, serving as policies, while effective in capturing multimodal action distributions, face two main challenges when applied to AHT. (1) While diffusion models excel at generating diverse actions through iterative denoising, they are primarily designed for distribution reconstruction and lack the predictive ability required for effectively supporting decision-making. This limitation raises challenges for applying diffusion models to AHT, where anticipating teammate intentions and making context-aware decisions are crucial. (2) Traditional diffusion models are typically built on MLP or UNet architectures, which lack the architectural flexibility to adapt to changing teammate behaviors. This presents a limitation for AHT, where real-

^{*}corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

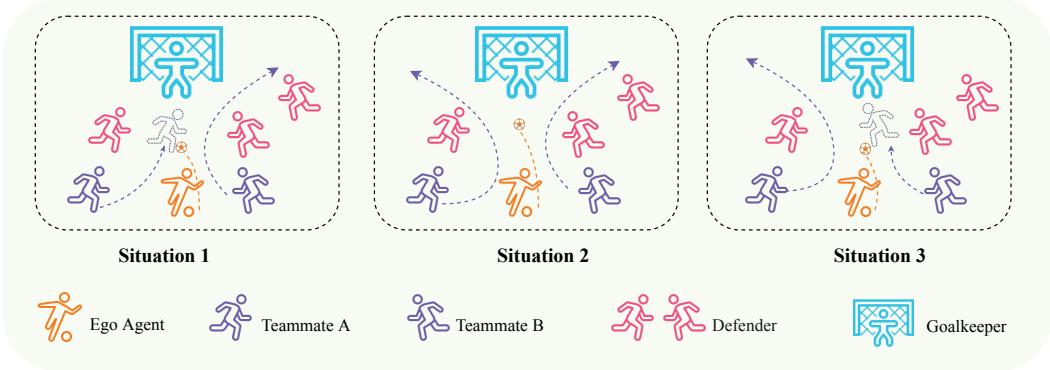


Figure 1: An example illustrating multiple cooperation patterns in AHT. Situation 1: The ego agent passes the ball to teammate A; Situation 2: The ego agent attempts a shot directly; Situation 3: The ego agent passes the ball to teammate B.

time adaptation to dynamic and unpredictable teammates is essential for effective collaboration.

To address these two limitations, we present PADiff, a diffusion-based framework for learning Predictive and Adaptive policies in AHT. The core innovations of PADiff are as follows. (1) To address the lack of predictive capability in traditional diffusion models, we design a **Predictive Guidance Block (PGB)** integrated into the denoising process. This module specifically leverages intermediate representations to predict teammates' cooperative targets and align action generation with long-term team objectives. (2) To address the limited adaptability of traditional diffusion architectures, we design an **Adaptive Feature Modulation Net (AFM-Net)**, which integrates two FiLM-like feature-wise modulation layers to scale and shift intermediate features based on the current state (Perez et al. 2018). To further improve the diffusion policies' stability and adaptability against uncertain teammates, AFM-Net integrates several designs such as Layer Normalization, Residual connections and Dropout regularization into the training process. Through this way, AFM-Net dynamically adjusts its internal representations to accommodate the frequently changing teammates' goals and behaviors.

The main contributions of our work can be summarized as follows.

1. We introduce PADiff, which is the first diffusion-based approach to the problem of AHT. PADiff enables the agent to learn multimodal cooperation patterns, leading to diverse collaboration with unknown teammates. This addresses the limitation of traditional RL-based methods which collapse to a single dominant behavior.
2. PADiff integrates two novel modules to enhance the diffusion policies' adaptability and prediction ability. First, the AFM-Net enables real-time adaptation to non-stationary teammates, ensuring robust performance in dynamic AHT scenarios. Second, the PGB effectively guides the denoising process during training by predicting teammates' cooperative targets.
3. We empirically validate PADiff in three classic collab-

orative environments. Experimental results demonstrate that PADiff consistently outperforms strong baselines in diverse scenarios, achieving an impressive average performance gain of 35.25% across all tested environments.

Related works

Ad Hoc Teamwork. The core challenge of AHT is to develop a policy that generalizes across diverse teammate behaviors and enables rapid adaptation during execution. Early approaches typically rely on predefined sets of teammate types, using type inference to switch among corresponding policies (e.g., PLASTIC (Barrett and Stone 2015)). While effective in constrained environments, these methods (Chen et al. 2020; Barrett and Stone 2015; Durugkar, Liebman, and Stone 2020; Mirsky et al. 2020) struggle when teammates exhibit complex or unmodeled behaviors. More recent works (Papoudakis, Christianos, and Albrecht 2021; Gu et al. 2021) improve adaptability by learning representations of teammates, allowing the agent to adjust its policy based on online interactions. These methods relax the assumption of fixed types and broaden the space of potential collaborations. However, most existing approaches (Barrett et al. 2017; Chen et al. 2020; Rahman et al. 2021; Zhang et al. 2025) are still built on RL frameworks, where the optimization objective focuses on maximizing expected return. This leads the learned policy to converge toward a single high-reward action, neglecting the diversity of possible cooperative modes under different teammate dynamics. While maximum entropy regularization ((Haarnoja et al. 2018)) mitigates this collapse by promoting exploratory behavior, it fundamentally fails to construct genuinely multimodal policies. Such limited policy expressiveness directly constrains generalization in AHT. In contrast, diffusion models (Zhu et al. 2023) have the ability to represent any normalizable distribution, with the potential to improve the AHT policy representation effectively.

Generative Models for Decision Making. Generative Models for Decision Making. Generative models have seen increasing use in decision-making tasks (Gozalo-Brizuela

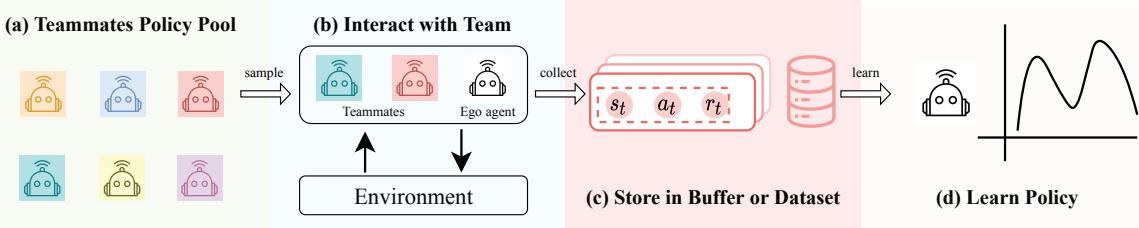


Figure 2: Overview of the AHT training pipeline. The ego agent learns to cooperate with diverse teammates by: (a) sampling from a heterogeneous teammate policy pool to simulate varied collaboration scenarios, (b) interacting within the environment to collect trajectories, (c) storing interaction data, and (d) continuously optimizing its policy based on the collected data.

and Garrido-Merchan 2023; Jo 2023; Brynjolfsson, Li, and Raymond 2025; Goodfellow et al. 2020). While GAN-based methods (Goodfellow et al. 2020) like GAIL (Ho and Ermon 2016) offer adversarial imitation, they often struggle with training instability and limited generalization. VAEs (Kingma, Welling et al. 2013; Mandlekar et al. 2020; Mees, Hermann, and Burgard 2022; Lynch et al. 2020) utilize latent representations for policy generation but face difficulties capturing fine-grained dynamics. More recently, Transformer-based models (Chen et al. 2021; Zheng, Zhang, and Grover 2022; Xie et al. 2023), notably Decision Transformer (Chen et al. 2021), have reframed policy learning as return-conditioned sequence modeling, achieving strong offline RL performance. However, most generative policy models remain unimodal due to architectural biases like Gaussian priors and autoregressive decoding. In contrast, diffusion-based policies (Wang, Hunt, and Zhou 2022) can learn multimodal action distributions. Leveraging this, we embed teammate-predictive guidance into a diffusion-based policy framework specifically for AHT.

Diffusion Models. Diffusion models (Sohl-Dickstein et al. 2015), first for general high-dimensional data, rapidly found success in image synthesis. Their strength here lies in U-Net backbones’ multi-scale feature aggregation and spatial inductive bias (Ho, Jain, and Abbeel 2020; Nichol and Dhariwal 2021; Vahdat, Kreis, and Kautz 2021). DiT (Peebles and Xie 2023) replaces the U-Net backbone with Vision Transformers, showing that attention mechanism improves sample quality by capturing long-range dependencies. Beyond vision, diffusion models have recently been adopted for decision-making in the RL field. Trajectory-level planners (Janner et al. 2022; Liang et al. 2023, 2024; Zhang et al. 2024) generate full state-action sequences for long-horizon reasoning, typically with U-Nets (Janner et al. 2022) or Transformer variants like DiT (Zhang et al. 2024). Action-level policies, by contrast, denoise one action at a time with lightweight MLPs (Wang, Hunt, and Zhou 2022; Kang et al. 2023; Hansen-Estruch et al. 2023), enabling real-time control. We follow the action-level paradigm but tackle the more demanding AHT setting. To this end, we introduce (i) a *Conditional Feature Modulation Network* that sustains robustness under dynamically changing teammate policies and (ii) a *Predictive Guidance Block* that injects teammate-aware goals into the denoising process.

Preliminaries

Ad Hoc Teamwork. We consider training an ego agent to collaborate with unknown teammates in cooperative multi-agent environments to achieve a shared goal. This problem is modeled as a Decentralized Markov Decision Process (Dec-MDP) extended with a teammate policy space: $\langle N, \mathcal{S}, \mathcal{A}, P, R, \Gamma, \mathcal{T} \rangle$, where $N = \{1, 2, \dots, n\}$ denotes the set of agents, \mathcal{S} is the global state space, and the joint action space is defined as $\mathcal{A} = \mathcal{A}^1 \times \dots \times \mathcal{A}^N$. Without loss of generality, we designate agent i as the ego agent and $-i$ as the set of teammates. The transition function $P : \mathcal{S} \times \mathcal{A} \rightarrow \Delta(\mathcal{S})$ governs the environment dynamics, and R is the shared reward function. Γ represents the joint policy space of all potential teammates. We define a trajectory $\tau = \{(\mathbf{s}_t, \mathbf{a}_t, r_t)\}_{t=0}^T$ and denote by \mathcal{T} the set of such episodic trajectories generated from interactions between the ego agent and sampled teammates.

As illustrated in Figure 2, the AHT training pipeline contains four stages: (a) Sample a subset of policies from a diverse teammate pool at the beginning of each episode; (b) Interact with the ego agent to collect full trajectories τ ; (c) Store trajectories in the offline dataset; (d) Optimize the ego policy $\pi_\theta^i(a^i | s; \mathcal{D})$ based on the collected data \mathcal{D} .

Diffusion Probabilistic Models. Diffusion models (DMs) (Sohl-Dickstein et al. 2015; Song and Ermon 2019) are a powerful type of generative models that learn to recover data samples from noise. This typically involves two phases: a forward noising process that gradually perturbs clean data into noise, and a reverse process that learns to reconstruct data from noise. Given a data sample $x_0 \sim p_{\text{data}}(x)$, the forward process defines a Markov chain: $p(x_k | x_{k-1}) = \mathcal{N}(x_k | \sqrt{\alpha_k}x_{k-1}, (1 - \alpha_k)I)$, where α_k controls the noise level at step k . The reverse process is parameterized by: $q_\theta(x_{k-1} | x_k) = \mathcal{N}(x_{k-1} | \mu_\theta(x_k, k), (1 - \alpha_k)I)$. Sampling starts from $x_K \sim \mathcal{N}(0, I)$ and iteratively denoises toward x_0 . The model is trained by minimizing the loss:

$$\mathcal{L}(\theta) = \mathbb{E}_{k, x_0, \epsilon} [\|\epsilon - \epsilon_\theta(x_k, k)\|^2], \quad \epsilon \sim \mathcal{N}(0, I). \quad (1)$$

Methods

Overview. Our framework PADiff leverages a diffusion-based approach tailored for cooperative decision-making in

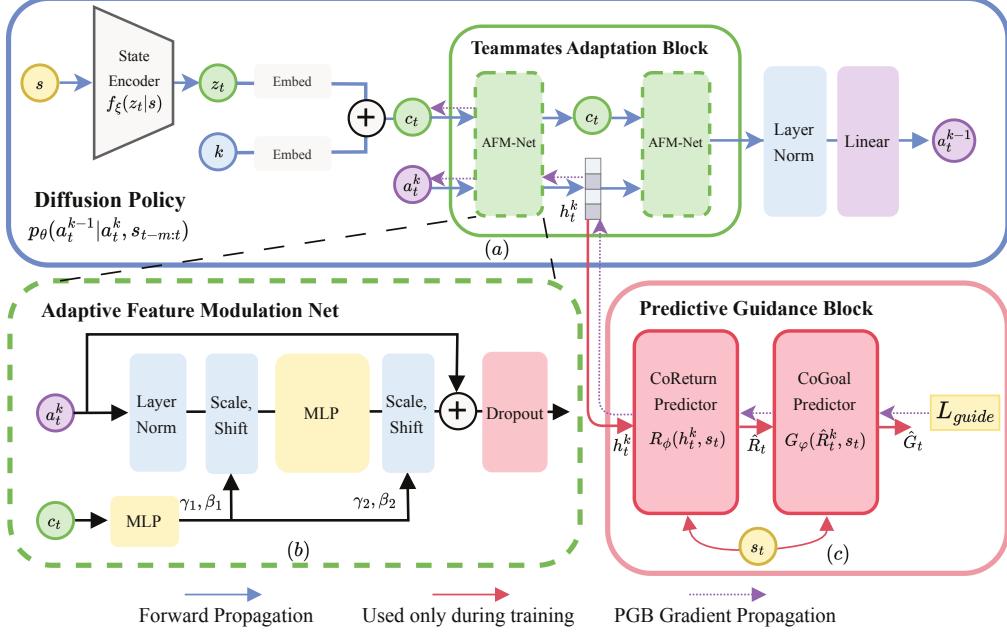


Figure 3: The overall architecture of PADiff. The subgraph (a) shows that PADiff represents the policy as a diffusion-based model and the State Encoder $f_\xi(z_t|s_{t-m:t})$ transforms states into latent representation capturing team dynamics. c_t represents the vector obtained by bit-wise addition of the embedding of latent representations z_t and the embedding of diffusion step k , which serves as the condition for diffusion. The subgraph (b) showcases our AFM-Net uses c_t representing teamwork context as the denoising condition to dynamically modulate the intermediate feature vectors to affect the actions generation. The subgraph (c) illustrates that PGB integrates the team-aware information into the action denoising process by predicting teammates’ intentions while training, through gradient propagation, ensuring that the ego agent can make such decisions that align with long-term team objectives while testing.

AHT environments. As illustrated in Figure 3, PADiff comprises three key components: (1) diffusion-based policy representation, (2) Teammates Adaptation Block, and (3) Predictive Guidance Block (PGB). Initially, the state encoder transforms the global state into a latent variable encoding the current teamwork context. The latent representation, together with action embedding, enters the Teammates Adaptation Block, consisting of multiple Adaptive Feature Modulation Nets (AFM-Net), which dynamically adjusts feature vectors via a FiLM-based conditioning mechanism, generating adaptive actions responsive to teammate dynamics. Simultaneously, the PGB predicts collaborative return and collaborative goal from intermediate features h_t^k of Teammates Adaptation Block, guiding the denoising process toward optimal team-aligned objectives by integrating predictive information via gradient signals.

During training, these three components jointly enable the policy model to have both predictive and adaptive ability. At inference, since the denoising module internalizes team-awareness, the PGB is no longer required, enabling efficient real-time adaptive decision-making. The procedure can be found in Alg1 and Alg2 of the Appendix. In the following, we detail each core component of PADiff.

Overall Diffusion Policy Representations

To capture agent’s potential multimodal behaviors required for effective ad hoc teamwork, unlocking its diverse cooperation modes with teammates, we model the ego agent’s policy as a conditional diffusion process. The iterative denoising framework naturally accommodates complex conditions, enabling us to inject compact teammates information at each step. We define the policy distribution over actions a_t given state s_t as the reverse diffusion denoising chain:

$$\pi_\theta(a_t|s_t) = p_\theta(a_t^{0:K}|s) = \mathcal{N}(a_t^K; 0, I) \prod_{k=K}^0 p_\theta(a_t^{k-1}|a_t^k, s), \quad (2)$$

As shown in figure 4, the solid lines show an inverse diffusion process where a noisy action a_t^K sampled from a Gaussian distribution progressively transforms into the final executable action a_t^0 through K diffusion steps while execution.

For discrete action spaces, following D3PM (Austin et al. 2021), the reverse diffusion is parameterized as:

$$a^{k-1}|a^k \sim p_\theta(a^{k-1}|s, a^k), \quad k = K, \dots, 1, \quad (3)$$

and the forward diffusion process can be represented as:

$$q(a_t^k|a_t^0) = \text{Cat}(a_t^k; p = a_t^0 \bar{Q} k). \quad (4)$$

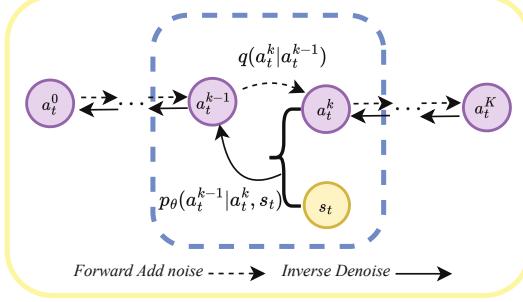


Figure 4: Forward adding noise process and inverse denoising process in PADiff.

$$q(\mathbf{a}_t^k | \mathbf{a}_t^{k-1}) = \text{Cat}(\mathbf{a}_t^k; \mathbf{p} = \mathbf{a}_t^{k-1} \mathbf{Q}_k) \quad (5)$$

where $\overline{\mathbf{Q}}_k = \mathbf{Q}_1 \cdot \mathbf{Q}_2 \dots \mathbf{Q}_k$ and the schedule type of transition matrices \mathbf{Q} follows the Uniform approach, which is a way of controlling the forward diffusion process and the learnable reverse diffusion process. $\text{Cat}(\mathbf{x}; \mathbf{p})$ is a categorical distribution over the one-hot row vector \mathbf{x} with probabilities given by the row vector \mathbf{p} (Austin et al. 2021).

During training, we sample k from $\{1, \dots, K\}$. As shown in figure 4, the dotted lines direction shows that the sample is prepared by starting with ground truth action \mathbf{a}_t^0 . We can compute \mathbf{a}^k and \mathbf{a}^{k-1} using equation 5. Then the model learns to model the entire distribution $p_\theta(\mathbf{a}_t^{k-1} | \mathbf{a}_t^k, s_t)$ by optimizing the variational lower bound. The loss function for training is as follows:

$$\begin{aligned} \mathcal{L}_{Diff} = \mathbb{E} [D_{KL} [q(\mathbf{a}^{k-1} | \mathbf{a}^k, \mathbf{a}^0) \| p_\theta(\mathbf{a}^{k-1} | \mathbf{a}^k, s, k)]] \\ - \log p_\theta(\mathbf{a}^0 | \mathbf{a}^1, s) \end{aligned} \quad (6)$$

where

$$\begin{aligned} q(\mathbf{a}_t^{k-1} | \mathbf{a}_t^k, \mathbf{a}_t^0) &= \frac{q(\mathbf{a}_t^k | \mathbf{a}_t^{k-1}, \mathbf{a}_t^0) q(\mathbf{a}_t^{k-1} | \mathbf{a}_t^0)}{q(\mathbf{a}_t^k | \mathbf{a}_t^0)} \\ &= \text{Cat} \left(\mathbf{a}_t^{k-1}; \mathbf{p} = \frac{\mathbf{a}_t^k \mathbf{Q}_k^\top \odot \mathbf{a}_t^0 \overline{\mathbf{Q}}_{k-1}}{\mathbf{a}_t^0 \overline{\mathbf{Q}}_k \mathbf{a}_t^k} \right). \end{aligned}$$

While our diffusion policy captures the full action distribution, standard diffusion models are insufficient for sophisticated ad hoc teamwork. They typically lack (1) efficient adaptation to dynamic teammates and (2) integrated teammate intent prediction. PADiff addresses these limitations with its novel Adaptive Feature Modulation Net and Predictive Guidance Block, which we will detail in the following.

State-Conditioned Teammates Adaptation

The denoising process $p_\theta(\mathbf{a}_t^k, s_t, k)$ uses a state encoder and the Teammates Adaptation Block. The state encoder compresses the state into a more informative latent variable. Then the latent variable, diffusion step and the action embedding are fed into the Teammates Adaptation Block, which is consisting of two AFM-Net, simultaneously to generate the denoised action.

State Encoder: To better identify the cooperation situation of the team and handle potential changes in uncertain teammates, we model the team’s cooperation state using a history window $s_{t-m:t}$ which provides richer context about interactions. This sequence is compressed into an informative latent space modeled as a multivariate Gaussian distribution (Gu et al. 2021). A sample from this latent distribution acts as the conditional input during the denoising stage.

$$(\mu_{z_t}, \sigma_{z_t}) = f_\xi(s_{t-m:t}), z_t \sim \mathcal{N}(\mu_{z_t}, \sigma_{z_t}) \quad (7)$$

Adaptive Feature Modulation Net (AFM-Net): Traditional denoising networks, such as UNet and MLP, lack necessary teammate-awareness, while Transformer-based networks like DiT (Peebles and Xie 2023) rely on attention mechanisms that introduce significant computational overhead, undesirable in fast-paced AHT settings. To overcome these limitations, we propose AFM-Net, an effective and robust denoising network that integrates the FiLM mechanism, residual connections, layer normalization and dropout regularization, making it both expressive and efficient without the need for attention. AFM-Net has three primary features: (1) Conditional feature modulation through FiLM-style (Perez et al. 2018) method, enhancing teammate adaptability, (2) Residual connections for stable training and robust representations and (3) Dropout regularization to improve generalization to unseen teammates.

Specifically, we generate γ_j (scale) and β_j (shift) parameters from the team context z_t using an MLP. These parameters dynamically modulate the intermediate feature vectors produced by Layer Norm Block and MLP Block. This integration enables the model to generate actions that adapt to the team. Additionally, residual connections and dropout regularization ensure stable learning and enhance the model’s ability to generalize to new, unseen teammates.

$$\gamma_1, \beta_1, \gamma_2, \beta_2 = \text{MLP}(\mathbf{z}_t + k) \quad (8)$$

$$\begin{aligned} \text{AFM}(\mathbf{x}, z_t, k) = \gamma_2 \cdot (\text{MLP}(\gamma_1 \cdot \text{LN}(\mathbf{x}) + \beta_1)) \\ + \beta_2 + \mathbf{x} \end{aligned} \quad (9)$$

where \cdot represents the element-wise dot product, \mathbf{x} represents the intermediate feature vector, LN represents the Layer Normalization module.

Predictive Denoising with Collaborative Goal Guidance

To address the challenge that existing diffusion-based policies lack the ability to predict teammates’ intentions and struggle to adapt to the behaviors of previously unseen teammates, we propose the **Predictive Guidance Block (PGB)**, a novel component integrated into the denoising phase of the PADiff framework, which can integrate predictive information into the denoising process.

Unlike conventional inference-time guidance methods, which often lack flexibility in adapting to real-time teammate dynamics, PGB leverages the intermediate representation generated during the denoising process to predict teammates’ intentions. By propagating gradients through the denoising process, it optimizes this representation, enabling it

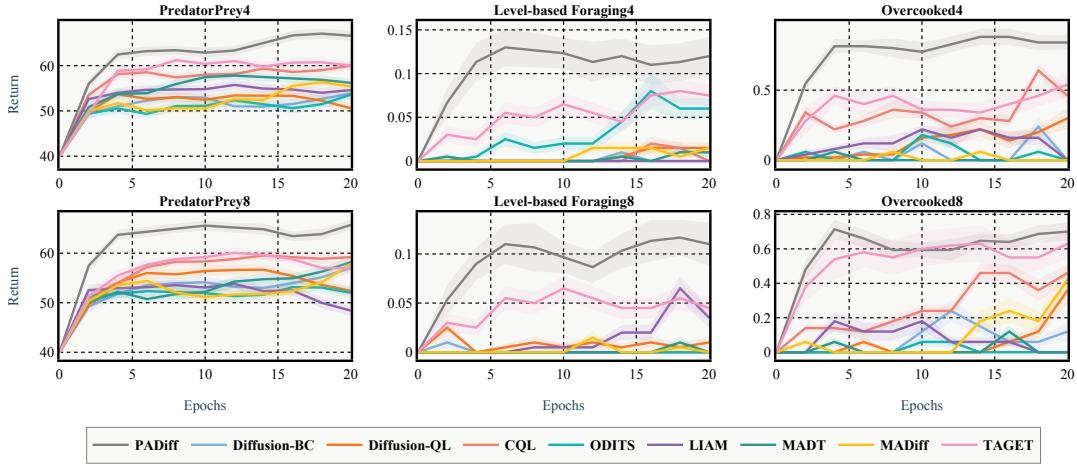


Figure 5: Average evaluation returns with 95% confidence interval during training.

to anticipate teammates’ behaviors. As a result, the denoising process becomes adaptive to changes in teammates’ actions. The design of the PGB contains two core predictive tasks: **Collaborative Return (CoReturn)** represents the expected cumulative future team reward, conditioned on the current state and intermediate feature from the denoising network. The prediction of CoReturn serves to guide the generation of actions towards better alignment with overall team objectives. And **Collaborative Goal (CoGoal)** stands for the predicted future state, which is served as a sub-goal to capture the intention of teammates. While model-base method like DreamerV3(Hafner et al. 2023) optimize policies through environment prediction, PGB similarly enhances decision-making via auxiliary prediction tasks, but fundamentally differs by tailoring predictions to real-time collaborative intent. Each predictive task contributes distinctively to enhanced robust cooperative decision-making.

During the training phase, we predict CoReturn R_t using the meaningful intermediate features h_t^k directly produced by AFM-Net and state s_t , and then use CoReturn along with s_t to predict CoGoal G_t . This sequential prediction process is designed to model the hierarchical nature of team coordination. The initial prediction of CoReturn provides an evaluation of the expected reward, which serves as a foundational signal for determining the next sub-goals, which are captured by CoGoal. The loss function is defined as:

$$L_{\text{CoReturn}} = \mathbb{E}_{\tau \sim \mathcal{D}} \left[\sum_{t=1}^T \|R_\phi(h_t^k, s_t) - R_t\|^2 \right], \quad (10)$$

$$L_{\text{CoGoal}} = \mathbb{E}_{\tau \sim \mathcal{D}} \left[-\frac{1}{N} \sum_{t=1}^T \sum_{i=1}^N \left(G_{t,i} \log \hat{G}_{t,i} + (1 - G_{t,i}) \log (1 - \hat{G}_{t,i}) \right) \right] \quad (11)$$

where $\hat{G}_t = G_\varphi(\hat{R}_t^k, s_t)$, $G_t = s_{t+n}$, $R_t = \sum_{t'=t}^T r_{t'}$. We perform one-hot encoding on future state s_{t+n} , where each dimension represents a binary indicator. We use BCE loss per dimension to handle binary classification, which aligns with the discrete nature of our environments.

The total loss is represented as:

$$L_{\text{total}} = L_{\text{Diffusion}} + \alpha L_{\text{CoReturn}} + \beta L_{\text{CoGoal}} \quad (12)$$

Formally, let $\nabla_{h_t^k} L_{\text{CoReturn}}$ and $\nabla_{h_t^k} L_{\text{CoGoal}}$ represent the gradient derived from each prediction objective. The total gradient update on h_t^k can be represented as:

$$\nabla_{h_t^k} L_{\text{total}} = \nabla_{h_t^k} L_{\text{Diffusion}} + \alpha \nabla_{h_t^k} L_{\text{CoReturn}} + \beta \nabla_{h_t^k} L_{\text{CoGoal}}, \quad (13)$$

where α, β are weighting hyperparameters. By minimizing prediction errors on CoReturn and CoGoal, intermediate features become predictive of beneficial cooperative targets, ensuring that the ego agent can respond effectively to teammates’ behaviors. During the inference phase, the Predictive Guidance Block is no longer required, as the denoise block AFM-Net has learned the capacity to produce intermediate features that are predictive of team-aware behavior, which enables the denoising block to embed team-awareness into its action generation and maintains effective coordination capacity without real-time guidance.

Experiments

Setup. To better evaluate the performance of the proposed model in Ad Hoc Teamwork scenarios, we employ three benchmark environments: **Predator-Prey (PP)**, **Level-Based Foraging (LBF)**, and **Overcooked**. Predator-Prey: In a grid-world environment, a team of predator agents must collaborate to capture a prey that actively attempts to evade them. Successful captures require tight coordination among predators. Level-Based Foraging: Agents operate in a grid world where they collect food items randomly distributed across the environment. Both agents and food items

are assigned discrete levels, and a food item can only be collected when the combined levels of participating agents meet or exceed the food's level. **Overcooked**: Agents must efficiently prepare and deliver dishes in a kitchen with a complex spatial layout. Effective cooperation with teammates and synchronized movements are crucial.

Collections of teammates policies. In AHT scenarios, we must ensure that the ego agent can adapt to previously unseen teammates. To build a diverse pool of teammate policies, we use the Soft-Value Diversity (SVD) method from the CSP(Ding et al. 2023) framework to collect teammate behaviors, instantiating four independent multi-agent populations in each environments. The training process alternates between the inner loop and the outer loop. In the inner loop, each team is trained in isolation, updating its policy to maximize its own performance. In the outer loop, we jointly update all teams to maximize the SVD objective, encouraging them to develop distinct value estimates over observation-action pairs, thus diversifying their behaviors. Specifically, three training populations are dedicated to interacting with the ego agent for collecting the training data, while a fourth population is reserved exclusively for assessing the ego agent's performance. Each training population comprises three unique policy checkpoints, whereas the testing population includes twelve checkpoints. These checkpoints correspond to distinct joint policies. By isolating the training and testing in this way, our protocol ensures a rigorous evaluation of the model's ability to adapt to previously unseen teammate strategies. We divided the 12 test policies into two groups, one containing 4 test strategies and the other containing 8. The cross-play matrices experiments reveal that cooperation scores for different teammate combinations are lower than for identical teammate combinations. This is due to strategy differences causing coordination issues, which reflects the diversity of test teammates strategies. Detailed information about cross-play experiments can be found in Appendix A.2.

Baselines. In our experiments, we compare two classic AHT methods: **LIAM** (Papoudakis, Christianos, and Albrecht 2021) and **ODITS** (Gu et al. 2021). To underscore the benefits of our method and its superior adaptability to AHT tasks, we further benchmark it against two diffusion-based models: **Diffusion-BC** and **Diffusion-QL** (Wang, Hunt, and Zhou 2022). We also compare with the offline reinforcement learning algorithm **CQL** (Kumar et al. 2020) and two AHT-focused multi-agent methods: **MAD-IFF** (Zhu et al. 2024), which aggregates all agents' observations into a global state sequence and uses an inverse dynamics model to predict the ego agent's next action; and **MADT** (Meng et al. 2021), which only trains the ego agent's actor network within the AHT settings. We also take the latest DT-based offline AHT method **TAGET** (Zhang et al. 2025) as baseline. These comparisons demonstrate that our method addresses the challenges in AHT scenarios more effectively and achieve significant performance.

Comparison with Baselines.

We trained our model for 20 epochs in the offline datasets we collected. We evaluate the agent against a held-out test pool

every two epochs,. Test pools consist of either 4 or 8 policies. For each environment, one group of policies is drawn at random and paired with the ego agent for evaluation. Returns are averaged over 50 trials, with shaded regions representing 95% confidence intervals using the standard normal distribution formula: $\bar{x} = \pm 1.96 \cdot \frac{\sigma}{\sqrt{n}}$, where σ is the sample standard deviation and $n = 50$. In every environment and for both pool sizes, our model consistently surpasses the baseline(cf. Fig 5 and Tab1 of Appendix). Additionally, hyperparameters analysis and detailed network architecture can be found in Fig4 and Fig5 of the Appendix.

Visualization of Multimodal Policy Distributions

To verify that PADiff captures multimodal behaviors and discovers different modes of cooperation, we set up a game scenario in the PP environment and input the same state into the policy multiple times. We found that the ego agent would exhibit different cooperative paths as shown in Fig1 of Appendix, proving that our policy actually fits a multimodal distribution, enabling diverse behaviors under the same state. Further analysis can be found in Appendix A.3.

Ablation Studies

To validate the necessity of our proposed AFM-Net, we replace AFM-Net in our framework with two alternative architectures: MLP and U-Net, and evaluate performance across all environments. As shown in the Figure 7, our model consistently outperforms both variants in all environments. These results strongly demonstrate that, in the context of AHT tasks, the design of AFM-Net offers significant advantages and greater applicability compared to a simple MLP and the image-oriented U-Net architecture.

To evaluate the contribution of each component of PGB module to overall framework performance, we conducted an ablation study by systematically removing individual modules and assessing performance, allowing us to quantify the impact of each component. As depicted in the Fig 6, the results show that the full model consistently outperforms all ablated variants. Both removing the CoGoal predictor and the CoReturn predictor results in substantial degradation, indicating that the ability to predict future team rewards and subgoal are crucial for enabling the ego agent to collaborate with unseen teammates. These findings strongly support the rationality for the design of PGB, demonstrating that each component plays a vital role in enhancing performance.

Conclusions

In this work, we introduced PADiff, a novel diffusion-based framework that employs diffusion models as the ego agent's policy for AHT. Through extensive experiments, we validated that PADiff successfully enhances collaboration in dynamic, unseen teammate scenarios by integrating predictive information into the denoising process and incorporating adaptive mechanisms for dynamic teammate adaptation. This paves the way for more autonomous and reliable multi-agent systems capable of seamless collaboration in open, unpredictable realworld environments.

Acknowledgments

This work is supported by National Natural Science Foundation of China (NSFC) under the Youth Science Fund Project (Grant Number: 62506133) and the Guangdong Basic and Applied Basic Research Foundation (Grant Number: 2025A1515010247).

References

Austin, J.; Johnson, D. D.; Ho, J.; Tarlow, D.; and Van Den Berg, R. 2021. Structured denoising diffusion models in discrete state-spaces. *Advances in neural information processing systems*, 34: 17981–17993.

Barrett, S.; Rosenfeld, A.; Kraus, S.; and Stone, P. 2017. Making friends on the fly: Cooperating with new teammates. *Artificial Intelligence*, 242: 132–171.

Barrett, S.; and Stone, P. 2015. Cooperating with unknown teammates in complex domains: A robot soccer case study of ad hoc teamwork. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29.

Brynjolfsson, E.; Li, D.; and Raymond, L. 2025. Generative AI at work. *The Quarterly Journal of Economics*, qjae044.

Chen, L.; Lu, K.; Rajeswaran, A.; Lee, K.; Grover, A.; Laskin, M.; Abbeel, P.; Srinivas, A.; and Mordatch, I. 2021. Decision transformer: Reinforcement learning via sequence modeling. *Advances in neural information processing systems*, 34: 15084–15097.

Chen, S.; Andrejczuk, E.; Cao, Z.; and Zhang, J. 2020. Ateam: Achieving the ad hoc teamwork by employing the attention mechanism. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 7095–7102.

Croitoru, F.-A.; Hondu, V.; Ionescu, R. T.; and Shah, M. 2023. Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9): 10850–10869.

Ding, H.; Jia, C.; Guan, C.; Chen, F.; Yuan, L.; Zhang, Z.; and Yu, Y. 2023. Coordination Scheme Probing for Generalizable Multi-Agent Reinforcement Learning.

Durugkar, I.; Liebman, E.; and Stone, P. 2020. Balancing individual preferences and shared objectives in multiagent reinforcement learning. International Joint Conference on Artificial Intelligence.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2020. Generative adversarial networks. *Communications of the ACM*, 63(11): 139–144.

Gozalo-Brizuela, R.; and Garrido-Merchan, E. C. 2023. ChatGPT is not all you need. A State of the Art Review of large Generative AI models. *arXiv preprint arXiv:2301.04655*.

Gu, P.; Zhao, M.; Hao, J.; and An, B. 2021. Online ad hoc teamwork under partial observability. In *International conference on learning representations*.

Haarnoja, T.; Zhou, A.; Abbeel, P.; and Levine, S. 2018. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *International conference on machine learning*, 1861–1870. Pmlr.

Hafner, D.; Pasukonis, J.; Ba, J.; and Lillicrap, T. 2023. Mastering diverse domains through world models. *arXiv preprint arXiv:2301.04104*.

Hansen-Estruch, P.; Kostrikov, I.; Janner, M.; Kuba, J. G.; and Levine, S. 2023. Idql: Implicit q-learning as an actor-critic method with diffusion policies. *arXiv preprint arXiv:2304.10573*.

Ho, J.; and Ermon, S. 2016. Generative adversarial imitation learning. *Advances in neural information processing systems*, 29.

Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.

Janner, M.; Du, Y.; Tenenbaum, J. B.; and Levine, S. 2022. Planning with diffusion for flexible behavior synthesis. *arXiv preprint arXiv:2205.09991*.

Jo, A. 2023. The promise and peril of generative AI. *Nature*, 614(1): 214–216.

Kang, B.; Ma, X.; Du, C.; Pang, T.; and Yan, S. 2023. Efficient diffusion policies for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 36: 67195–67212.

Kingma, D. P.; Welling, M.; et al. 2013. Auto-encoding variational bayes.

Kumar, A.; Zhou, A.; Tucker, G.; and Levine, S. 2020. Conservative q-learning for offline reinforcement learning. *Advances in Neural Information Processing Systems*, 33: 1179–1191.

Liang, Z.; Mu, Y.; Ding, M.; Ni, F.; Tomizuka, M.; and Luo, P. 2023. Adaptdiffuser: Diffusion models as adaptive self-evolving planners. *arXiv preprint arXiv:2302.01877*.

Liang, Z.; Mu, Y.; Ma, H.; Tomizuka, M.; Ding, M.; and Luo, P. 2024. Skilldiffuser: Interpretable hierarchical planning via skill abstractions in diffusion-based task execution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16467–16476.

Lynch, C.; Khansari, M.; Xiao, T.; Kumar, V.; Tompson, J.; Levine, S.; and Sermanet, P. 2020. Learning latent plans from play. In *Conference on robot learning*, 1113–1132. Pmlr.

Mandlekar, A.; Xu, D.; Martín-Martín, R.; Savarese, S.; and Fei-Fei, L. 2020. Learning to generalize across long-horizon tasks from human demonstrations. *arXiv preprint arXiv:2003.06085*.

Mees, O.; Hermann, L.; and Burgard, W. 2022. What matters in language conditioned robotic imitation learning over unstructured data. *IEEE Robotics and Automation Letters*, 7(4): 11205–11212.

Meng, L.; Wen, M.; Yang, Y.; Le, C.; Li, X.; Zhang, W.; Wen, Y.; Zhang, H.; Wang, J.; and Xu, B. 2021. Offline pre-trained multi-agent decision transformer: One big sequence model tackles all smac tasks. *arXiv preprint arXiv:2112.02845*.

Mirsky, R.; Macke, W.; Wang, A.; Yedidsion, H.; and Stone, P. 2020. A penny for your thoughts: The value of communication in ad hoc teamwork. International Joint Conference on Artificial Intelligence.

Nichol, A. Q.; and Dhariwal, P. 2021. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, 8162–8171. PMLR.

Papoudakis, G.; Christianos, F.; and Albrecht, S. 2021. Agent modelling under partial observability for deep reinforcement learning. *Advances in Neural Information Processing Systems*, 34: 19210–19222.

Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4195–4205.

Perez, E.; Strub, F.; De Vries, H.; Dumoulin, V.; and Courville, A. 2018. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32.

Rahman, M. A.; Hopner, N.; Christianos, F.; and Albrecht, S. V. 2021. Towards open ad hoc teamwork using graph-based policy learning. In *International conference on machine learning*, 8776–8786. PMLR.

Ravula, M. C. R. 2019. *Ad-hoc teamwork with behavior-switching agents*. Ph.D. thesis.

Sohl-Dickstein, J.; Weiss, E.; Maheswaranathan, N.; and Ganguli, S. 2015. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, 2256–2265. pmlr.

Song, Y.; and Ermon, S. 2019. Generative modeling by estimating gradients of the data distribution. *Advances in neural information processing systems*, 32.

Stone, P.; Kaminka, G.; Kraus, S.; and Rosenschein, J. 2010. Ad hoc autonomous agent teams: Collaboration without pre-coordination. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 24, 1504–1509.

Teng, S.; Hu, X.; Deng, P.; Li, B.; Li, Y.; Ai, Y.; Yang, D.; Li, L.; Xuanyuan, Z.; Zhu, F.; et al. 2023. Motion planning for autonomous driving: The state of the art and future perspectives. *IEEE Transactions on Intelligent Vehicles*, 8(6): 3692–3711.

Vahdat, A.; Kreis, K.; and Kautz, J. 2021. Score-based generative modeling in latent space. *Advances in neural information processing systems*, 34: 11287–11302.

Wang, Z.; Hunt, J. J.; and Zhou, M. 2022. Diffusion policies as an expressive policy class for offline reinforcement learning. *arXiv preprint arXiv:2208.06193*.

Xie, Z.; Lin, Z.; Ye, D.; Fu, Q.; Wei, Y.; and Li, S. 2023. Future-conditioned unsupervised pretraining for decision transformer. In *International Conference on Machine Learning*, 38187–38203. PMLR.

Yuan, L.; Li, L.; Zhang, Z.; Chen, F.; Zhang, T.; Guan, C.; Yu, Y.; and Zhou, Z.-H. 2023a. Learning to Coordinate with Anyone. In *Proceedings of the Fifth International Conference on Distributed Artificial Intelligence*, 1–9.

Yuan, L.; Zhang, Z.; Li, L.; Guan, C.; and Yu, Y. 2023b. A survey of progress on cooperative multi-agent reinforcement learning in open environment. *arXiv preprint arXiv:2312.01058*.

Yucesoy, E.; Balcik, B.; and Coban, E. 2025. The role of drones in disaster response: A literature review of operations research applications. *International Transactions in Operational Research*, 32(2): 545–589.

Zhang, M.; Cai, Z.; Pan, L.; Hong, F.; Guo, X.; Yang, L.; and Liu, Z. 2024. Motiondiffuse: Text-driven human motion generation with diffusion model. *IEEE transactions on pattern analysis and machine intelligence*, 46(6): 4115–4128.

Zhang, X.; Chan, H.; Ye, D.; Cai, Y.; and Zhao, M. 2025. Ad Hoc Teamwork via Offline Goal-Based Decision Transformers. In *Forty-second International Conference on Machine Learning*.

Zheng, Q.; Zhang, A.; and Grover, A. 2022. Online decision transformer. In *international conference on machine learning*, 27042–27059. PMLR.

Zhu, Z.; Liu, M.; Mao, L.; Kang, B.; Xu, M.; Yu, Y.; Ermon, S.; and Zhang, W. 2024. Madiff: Offline multi-agent learning with diffusion models. *Advances in Neural Information Processing Systems*, 37: 4177–4206.

Zhu, Z.; Zhao, H.; He, H.; Zhong, Y.; Zhang, S.; Guo, H.; Chen, T.; and Zhang, W. 2023. Diffusion models for reinforcement learning: A survey. *arXiv preprint arXiv:2311.01223*.