

000 001 002 003 004 005 DIFFUSION-BASED BEHAVIOR CLONING IN MULTI- 006 AGENT GAMES VIA DYNAMIC GUIDANCE 007 008 009

010 **Anonymous authors**
011 Paper under double-blind review
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026

ABSTRACT

027 In offline multi-agent imitation learning, agents are constrained to learn from static
028 datasets without interaction, which poses challenges in generalizing across diverse
029 behaviors. Behavior Cloning (BC), a widely used approach, models conditional
030 actions from local observations but lacks robustness under behavioral variability.
031 Recent diffusion-based policies have been introduced to capture diverse action
032 distributions. However, in multi-agent environments, their iterative denoising pro-
033 cess can accumulate errors in interactive settings, degrading performance under
034 shifting opponent behaviors. To address these challenges, we propose Diffusion
035 Dynamic Guidance Imitation Learning (DDGIL), a diffusion-based framework
036 built on classifier-free guidance (CFG), which balances conditional and uncondi-
037 tional denoising predictions. Unlike prior methods with fixed weighting, DDGIL
038 introduces a dynamic guidance mechanism that adaptively adjusts the weight at
039 each denoising step, enhancing stability across different agent strategies. Empir-
040 ical evaluations on competitive and cooperative benchmarks show that DDGIL
041 achieves reliable performance. In high-fidelity sports simulations, it reproduces
042 action strategies that closely resemble expert demonstrations while maintaining
043 robustness against diverse opponents.
044

1 INTRODUCTION

045 Multi-Agent Reinforcement Learning (MARL) (Chongjie Zhang, 2010; Tabish Rashid, 2018; De-
046 heng Ye, 2020; Jiechuan Jiang, 2023) has focused on capturing inter-agent dependencies, and effective
047 training typically relies on reward signals obtained through repeated interactions with the environment.
048 However, in many real-world domains such as sports analytics, dense rewards are not available, and
049 designing explicit reward functions in sparse-reward settings is often unreliable (Jiexin Xie, 2019).
050 Moreover, online interaction is often infeasible due to cost or data collection constraints, leaving only
051 historical trajectories consisting of states, actions, and outcomes. These limitations motivate offline
052 imitation learning, where policies are derived solely from demonstrations without access to reward
053 signals or additional environment interaction.

054 In the offline setting, imitation learning methods, particularly behavior cloning (BC), are trained on
055 data collected against a fixed opponent, which constrains the learned policy to specific interaction
056 patterns. When evaluated against opponents with different behaviors, such as weaker or stylistically
057 distinct agents, these policies often exhibit unstable performance due to ineffective actions or mistimed
058 responses.

059 Recent studies have explored diffusion models for imitation learning, leveraging their ability to
060 represent complex action distributions (Cheng Chi, 2023) in an attempt to address the limitations of
061 BC. Existing approaches can be roughly divided into three categories. The first, often referred to as
062 Diffusion Policy (DP) (Tim Pearce, 2023; Zhendong Wang, 2023), treats the diffusion model as the
063 policy by conditioning primarily on the state and generating actions through multi-step denoising. As
064 illustrated in the lower part of Figure 1, the denoising process in this class of methods typically relies
065 solely on the observed condition, which fixes the denoising direction across steps. This rigidity makes
066 the process prone to error accumulation, particularly in interactive environments where observations
067 evolve with the behaviors of other agents, resulting in degraded performance. The second predicts
068 future states through diffusion and recovers actions using an inverse dynamics model (Michael Janner,
069 2022; Anurag Ajay, 2023). While effective in static settings, this approach can degrade in interactive
070 environments.

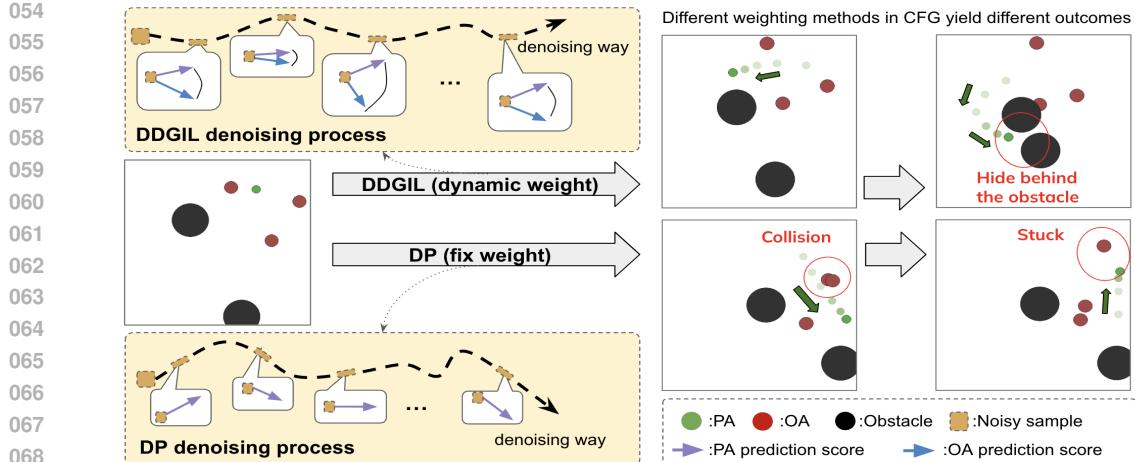


Figure 1: Tag task in the Multi-Agent Particle Environments (MPE), where three opponent agents (OA) pursue a primary agent (PA). We compare Diffusion Policy (DP, fixed weight) with DDGIL (dynamic weight). Under diverse opponent strategies, DP exhibits unstable behaviors, whereas DDGIL adapts and attains more stable performance.

domains when predicted states diverge from actual interactions. The third estimates conditional policy distributions for behavior cloning (Shang-Fu Chen, 2024), where the diffusion model provides auxiliary guidance during training. This additional signal can improve inference performance, yet the resulting policy remains essentially BC and continues to suffer from the same limitations when faced with diverse opponents. These methods indicate that diffusion-based approaches in offline imitation learning inherit key vulnerabilities of BC, highlighting the central challenge in multi-agent environments: achieving robustness under opponent variability.

Motivated by these challenges, we propose Diffusion Dynamic Guidance Imitation Learning (DDGIL), a diffusion-based imitation learning framework for offline multi-agent settings. DDGIL extends diffusion policies by introducing a minimal modification to the classifier-free guidance mechanism. As illustrated in the upper part of Figure 1, although our method employs the same denoising process as DP, the key difference lies in the use of a dynamic adjustment mechanism rather than a fixed weight. This dynamic design leverages a confidence signal to adapt to variations in the behavior of other agents, allowing the policy to remain flexible in interactive environments. Despite its simplicity, this modification improves stability and yields consistently better performance than fixed-weight diffusion policies and other baselines in both competitive and cooperative tasks.

2 RELATED WORK

Multi-Agent Imitation Learning (MAIL). Multi-Agent Imitation Learning (MAIL) has largely focused on online interaction with the environment during training (Yu et al., 2019; Nathaniel Haynam, 2025; Zare et al., 2024). Adversarial extensions such as MAGAIL (Jiaming Song, 2018) and CoDAIL (Minghuan Liu, 2020) adapt GAIL frameworks to the multi-agent setting, where a discriminator provides feedback to shape coordinated policies through interaction. While effective in capturing coordination strategies, these methods fundamentally rely on online rollouts, limiting their applicability to offline scenarios. STRIL (Shiqi Lei, 2025) represents one of the few offline approaches, filtering low-quality trajectories using strategy representations. However, it primarily focuses on modeling the agent’s own behavioral heterogeneity and does not explicitly account for opponent variability. In our work, we address this gap by proposing an offline MAIL framework that adapts robustly to diverse opponent strategies.

Diffusion Methods for Policy. Diffusion models (Jonathan Ho, 2020) have gained significant attention for their ability to iteratively denoise Gaussian noise and generate high-quality samples. In decision-making tasks, they have been applied to imitation and reinforcement learning by treating the diffusion model as the policy itself (Anurag Ajay, 2023; Michael Janner, 2022), generating

108 actions or trajectories that capture multi-modal distributions. While effective in single-agent and
 109 robotic control tasks, their application to multi-agent environments remains limited. Recent work,
 110 such as MaDiff (Zhengbang Zhu, 2024), presents an offline multi-agent RL framework that uses
 111 attention-based diffusion to model coordinated behavior.

112 A second line of work incorporates diffusion as an auxiliary component. DiffAIL (Bingzheng Wang,
 113 2024) augments adversarial imitation learning by adding a diffusion-based loss to the discriminator,
 114 improving its ability to distinguish expert and policy distributions. DBC (Shang-Fu Chen, 2024)
 115 is another representative offline method, which augments standard behavior cloning by combining
 116 a policy learning objective with an additional diffusion-based joint modeling loss. In both cases,
 117 diffusion serves as a supporting mechanism rather than the policy itself. In contrast, our work
 118 proposes to use diffusion models directly as policies, enabling stable and robust training in offline
 119 multi-agent imitation learning.

120

121 3 PRELIMINARIES

123 3.1 BEHAVIOR CLONING

125 Imitation learning (IL) aims to learn policies by replicating expert demonstrations without relying
 126 on explicit reward signals. In the multi-agent setting, we consider K agents, where each agent
 127 $i \in \{1, \dots, K\}$ receives an observation $s^{(i)}$ and selects an action $a^{(i)}$. We extract $(s^{(i)}, a^{(i)})$ pairs
 128 from a dataset \mathcal{D} . A common offline approach is behavior cloning (BC), which fits a policy by
 129 maximizing the likelihood of expert actions given the observations:

$$130 \pi^{(i)} = \arg \max_{\pi^{(i)}} \mathbb{E}_{(s_t^{(i)}, a_t^{(i)}) \sim \mathcal{D}} \left[\log \pi(a_t^{(i)} | s_t^{(i)}) \right] \quad (1)$$

132 where $\pi^{(i)}$ denotes the policy of agent i . While effective for replicating individual behaviors, BC
 133 conditions on local observations and does not explicitly account for inter-agent coupling, which may
 134 yield suboptimal coordination in multi-agent settings.

136 3.2 DIFFUSION PROBABILISTIC MODELS

138 Diffusion models (DMs), particularly denoising diffusion probabilistic models (DDPM) (Jonathan Ho,
 139 2020; Alexander Quinn Nichol, 2021), are generative models that learn to represent complex data
 140 distributions by gradually transforming Gaussian noise into structured samples through a multi-step
 141 process. We adopt this framework for action generation, allowing for high-dimensional policy
 142 modeling.

143

144 Decision Making. In our setting, the action $a \in \mathbb{R}^d$ is the generation target, conditioned on the current
 145 observation s . The forward process assumes access to state-action pairs (s, a) and adds Gaussian
 146 noise to a to produce a noisy sample x_t , defined as: $q(x_t | a) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t}a, (1 - \bar{\alpha}_t)I)$, $\bar{\alpha}_t =$
 147 $\prod_{s=1}^t \alpha_s$, where $\alpha_t \in (0, 1)$ is the predefined noise schedule. To recover the original action from
 148 the noisy input, a neural network ϵ_θ is trained to predict the added noise $\epsilon \sim \mathcal{N}(0, I)$. The noised
 149 input is computed as $x_t = \sqrt{\bar{\alpha}_t}a + \sqrt{1 - \bar{\alpha}_t}\epsilon$, with $t \sim \text{Uniform}(1, T)$. The model minimizes the
 150 following objective:

$$151 \mathcal{L}_{\text{DM}} = \mathbb{E}_{(s, a) \sim \mathcal{D}, t \sim \mathcal{U}(1, T), \epsilon \sim \mathcal{N}(0, I)} \left[\|\epsilon_\theta(x_t, t | s) - \epsilon\|^2 \right] \quad (2)$$

153

154 At inference time, the denoising process is performed iteratively. At each reverse step t , the model
 155 predicts the noise $\epsilon_\theta(x_t, t | s)$ and uses it to compute a denoised mean, which defines the mean of a
 156 Gaussian distribution used to calculate the noisy sample at the previous timestep x_{t-1} :

$$157 \mu_\theta(x_t, t | s) = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t | s) \right), \quad x_{t-1} = \mu_\theta(x_t, t | s) + \sigma_t \cdot z, \quad z \sim N(0, I) \quad (3)$$

160

161 where $\sigma_t^2 = \tilde{\beta}_t = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} (1 - \alpha_t)$, and $\sigma_t = \sqrt{\tilde{\beta}_t}$. Our imitation learning framework builds on this
 diffusion-based policy structure and extends it to multi-agent settings.

162 **Classifier-Free Guidance.** To improve conditional control during sampling, we apply classifier-free
 163 guidance (CFG) (Jonathan Ho, 2022). This method mixes the model’s unconditional and conditional
 164 predictions during denoising:

$$\epsilon_{\text{CFG}} = w \cdot \epsilon_{\theta}(x_t, t \mid s) + (1 - w) \cdot \epsilon_{\theta}(x_t, t) \quad (4)$$

167 Here, $\epsilon_{\theta}(x_t \mid s)$ is the conditionally guided prediction, while $\epsilon_{\theta}(x_t)$ is the unconditional estimate, and
 168 $w \in [0, 1]$ is a scalar that controls the guidance strength. Throughout this paper, we adopt a convex-
 169 combination form of classifier-free guidance, mixing conditional and unconditional predictions with
 170 a weight $w \in [0, 1]$. A larger w increases the adherence to the conditioning signal, while a smaller w
 171 yields more diverse samples. At inference time, ϵ_{CFG} is directly substituted into the denoising update
 172 (Eq. 3) to compute the reverse steps from x_t to x_{t-1} . Compared with classifier-based guidance,
 173 CFG avoids the difficulty of training a separate classifier under noisy inputs and has been widely
 174 observed to achieve stronger conditional fidelity and overall performance (Xi Wang & Kalogeiton,
 175 2024; Chung et al., 2025).

177 4 METHODS

179 Our objective is to address the instability of behavior cloning in multi-agent interaction, particu-
 180 larly when agents exhibit diverse or previously unseen strategies. To this end, we propose Diffu-
 181 sion Dynamic Guidance Imitation Learning (DDGIL), a diffusion-based framework that combines
 182 opponent-aware score prediction with a dynamic guidance mechanism during inference. Inspired by
 183 classifier-free guidance in diffusion-based image generation (Xi Wang & Kalogeiton, 2024; Chung
 184 et al., 2025), DDGIL replaces the fixed guidance weight with a per-step adjustment derived from the
 185 confidence of conditional predictions. This mechanism reduces the limitations of fixed weighting,
 186 enabling the policy to follow conditional signals when reliable and to adapt when they are uncertain.
 187 The overall architecture is shown in Figure 2.

188 4.1 PRIMARY AGENT DIFFUSION POLICY

190 Our method is based on offline imitation learning with diffusion models in a multi-agent setting with
 191 K agents. During training, we instantiate a separate diffusion model for each agent and designate
 192 one as the *primary agent*. The remaining agents are referred to as *opponent agents*, where the term
 193 “opponent” simply denotes agents other than the primary one: they act as adversaries in competitive
 194 tasks and as partners in cooperative tasks. **The choice of the primary agent or opponent agent is**
 195 **user-specified during inference and does not depend on the agent’s role.**

196 To train the primary agent policy, we sample state-action pairs (s_d^G, a_d^G) from the dataset \mathcal{D} , where
 197 $d \in 1, \dots, |\mathcal{D}| \times H$ indexes the pair and G denotes the primary agent. Here, $|\mathcal{D}|$ is the number of
 198 trajectories and H their length, yielding a total of $|\mathcal{D}| \times H$ state-action pairs.

199 The observation s_d^G serves as the conditional input, and Gaussian noise is added to the action a_d^G for
 200 diffusion model training. Following the denoising diffusion framework, we construct the input as:

$$x_t^G = \sqrt{\bar{\alpha}_t} a_d^G + \sqrt{1 - \bar{\alpha}_t} \epsilon, \quad \hat{\epsilon}_G = \epsilon_{\theta}^G(x_t^G, t \mid s_d^G) \quad (5)$$

203 where $\epsilon \sim \mathcal{N}(0, I)$, $t \sim \text{Uniform}(1, T)$, and train a diffusion model ϵ_{θ}^G to predict the noise. The
 204 training objective is defined as:

$$\mathcal{L}_{\text{DM}(G)} = \mathbb{E}_{(s_d^G, a_d^G) \sim \mathcal{D}, t \sim U(1, T), \epsilon \sim \mathcal{N}(0, I)} \left[\|\epsilon - \hat{\epsilon}_G\|^2 \right] \quad (6)$$

208 This objective is equivalent to denoising score matching and promotes the original policy distribution
 209 of the primary agent as observed in the dataset. Therefore, this section focuses on modeling the
 210 behavior of the primary agent, while the modeling of the remaining opponent agents will be described
 211 in the next Section 4.2.

213 4.2 OPPONENT-AWARE DIFFUSION MODELING

215 We consider K agents in total and write $k = K - 1$ for the number of opponents, with indices
 $i \in \{1, \dots, k\}$. For each opponent i , we train a separate diffusion model using the same procedure

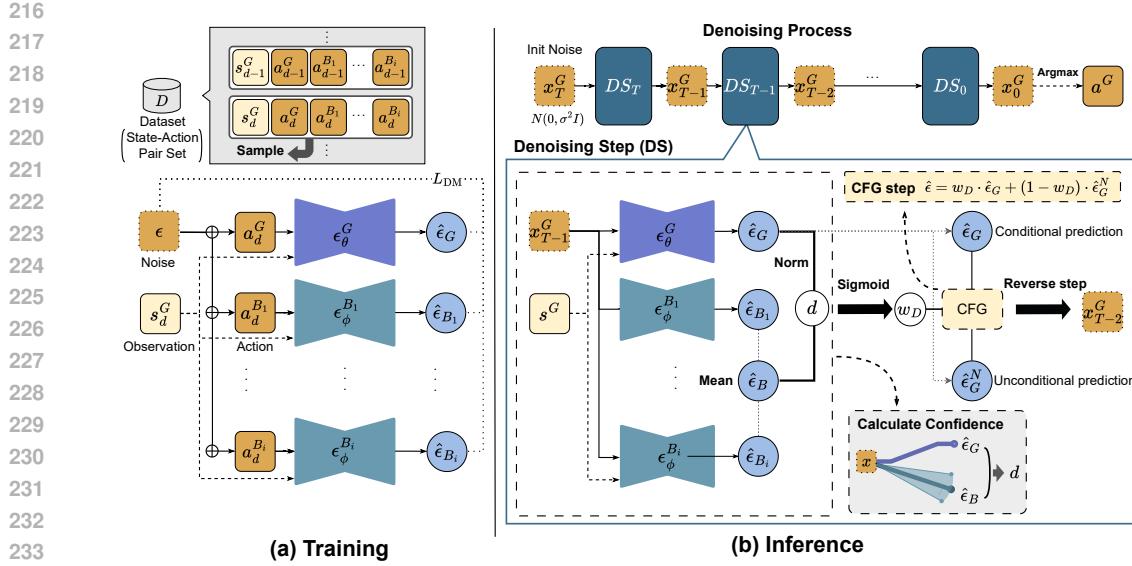


Figure 2: Overview of the **DDGIL** architecture: (a) the training pipeline, where separate diffusion-based policy models are trained for the main agent and other agents (opponents or collaborators); (b) the inference-time dynamic guidance mechanism, which computes weights at each denoising step based on the conditional scores from the primary and opponent agents, enabling adaptive policy adjustment based on the current interaction context.

as for the primary agent in Section 4.1. Concretely, we add Gaussian noise to the opponent’s action $a_d^{B_i}$ from the same dataset pair $(s_d^G, a_d^G, a_d^{B_i}, \dots)$, reuse the noise sample ϵ and diffusion step t , and condition the model on the primary agent’s observation s_d^G .

The noisy input and target prediction for each opponent are defined similarly to Eq. 5:

$$x_t^{B_i} = \sqrt{\bar{\alpha}_t} a_d^{B_i} + \sqrt{1 - \bar{\alpha}_t} \epsilon, \quad \hat{\epsilon}_{B_i} = \epsilon_{\phi}^{B_i}(x_t^{B_i}, t \mid s_d^G) \quad (7)$$

where $\epsilon_{\phi}^{B_i}$ predicts the diffusion noise for opponent B_i , conditioned on the shared context s_d^G . The conditioning on s_d^G provides a shared reference for aligning denoising trajectories and does not imply that opponents rely on the primary agent’s observation.

The overall training loss combines the reconstruction losses of the primary agent and all opponents, as shown in Figure 2(a). Since all models receive the same noise input ϵ , the ideal case is that both the primary agent and its opponents predict the same noise. **Sharing the same noise aligns the denoising trajectories across agents so that their disagreement reflects behavioral differences rather than randomness.** To encourage this, we add a latent consistency regularization term that minimizes the difference between their predicted noise, controlled by a coefficient c :

$$\mathcal{L}_{\text{DM(Joint)}} = \mathbb{E}_{(s_d^G, a_d^G, \{a_d^{B_i}\}_{i=1}^{|k|}) \sim \mathcal{D}} \left[\|\epsilon - \hat{\epsilon}_G\|^2 + \sum_{i=1}^{|k|} \|\epsilon - \hat{\epsilon}_{B_i}\|^2 + c \sum_{i=1}^{|k|} \|\hat{\epsilon}_G - \hat{\epsilon}_{B_i}\|^2 \right]. \quad (8)$$

This consistency term stabilizes the denoising dynamics under shared noise and is not intended to reduce or align behavioral differences between agents.

Recent studies show that diffusion models can reflect differences in noise through their denoising behavior (Bingzheng Wang, 2024; Mark S. Graham, 2023; Luping Liu, 2022; Yunshu Wu, 2024). Motivated by this property, we train a separate diffusion model for each opponent to capture its behavioral tendencies and decision patterns. During inference, these models do not directly generate actions; instead, their responses to shared noise inputs serve as guidance signals for the primary agent’s policy. In Section 4.3, we describe how to integrate these models into an opponent-aware mechanism.

270 4.3 DYNAMIC WEIGHT ADJUSTMENT DURING INFERENCE
271272 Each denoising step in the diffusion process updates a latent vector x_t , ultimately producing the
273 final action x_0 . We adopt the CFG framework, interpolating between conditional and unconditional
274 predictions, and replace Eq. 4 with

275
$$\hat{\epsilon} = w_D \cdot \hat{\epsilon}_G + (1 - w_D) \cdot \hat{\epsilon}_G^N \quad (9)$$

276

277 where $\hat{\epsilon}_G$ and $\hat{\epsilon}_G^N$ denote conditional and unconditional predictions, respectively. Unlike the standard
278 CFG setting that uses a fixed constant, we introduce a dynamic weight w_D that adapts at each
279 denoising step. This mechanism enables the policy to adjust its reliance on conditional information
280 when opponent behaviors render such predictions less reliable. To compute the weight w_D , we
281 measure the discrepancy between the agent’s conditional prediction and the noise predicted by
282 auxiliary opponent models. **These opponent models do not perform decision making or strategic
283 reasoning. They provide contextual score predictions that supplement the conditional estimate and
284 are incorporated through the following measure.**

285
$$d = \|\hat{\epsilon}_B - \hat{\epsilon}_G\|_2 \quad , \hat{\epsilon}_B = \mathbb{E}[\hat{\epsilon}_{B_i}] \quad , w_D = \sigma(d) = \frac{1}{1 + e^{-d}} \quad (10)$$

286

288 where $\hat{\epsilon}_B$ is the averaged prediction from the opponent models in the same denoising step. The value
289 d serves as a measure of confidence. Note that $d \geq 0$ guarantees $w_D \in [0.5, 1]$, so every update
290 remains a convex combination of $\hat{\epsilon}_G$ and $\hat{\epsilon}_G^N$. This lower bound ensures the update never favors the
291 unconditional branch, keeping conditioning as the default.292 The confidence d has two typical cases. When d is small, the two predictors are already aligned; in this
293 regime, increasing the weight of $\hat{\epsilon}_G$ contributes little useful signal and may amplify noise. Therefore,
294 the rule maintains a more conservative update close to $\hat{\epsilon}_G^N$, avoiding unnecessary fluctuation. When d
295 is large, the discrepancy indicates that the step carries important predictive information, and increasing
296 w_D allows the conditional term to contribute more strongly so that this information is preserved.297 Because w_D is recomputed at every denoising step, the mechanism adapts naturally along the
298 diffusion trajectory, where early steps usually provide weaker cues and later steps provide stronger
299 ones. This dynamic adjustment improves resilience against diverse and shifting opponent strategies,
300 avoiding the rigidity of a fixed coefficient. Section 5.3, 5.4, and 5.5 compare against fixed-weight
301 baselines, and Appendix A further analyzes the use of d as confidence and the properties of the
302 weighting rule.303 5 EXPERIMENTS
304306 To evaluate the effectiveness and generality of our proposed method **DDGIL**, we evaluate on
307 continuous, discrete, and combinatorial domains and compare against offline imitation baselines. Our
308 experiments are designed to address the following research questions:309

- **RQ1:** How well does DDGIL imitate expert behavior compared to baseline diffusion and
310 non-diffusion methods?
- **RQ2:** Can DDGIL adapt smoothly to opponents of varying strengths, with stable behavioral
311 and reward transitions?
- **RQ3:** How well does DDGIL perform in real-world or high-fidelity scenarios requiring
312 strategic generalization?

317 5.1 DATA COLLECTION AND EVALUATION SETUP

318 **Data Collection.** For each environment, we train a task-specific reinforcement learning policy and
319 use the best checkpoint as the expert policy O_{expert} to collect an offline dataset \mathcal{D} containing both
320 successful and failed trajectories.322 We additionally prepare two evaluation opponents: a mid-training policy (O_{medium}) and an early-stage,
323 near-random policy (O_{weak}). These checkpoints are used solely for evaluation to test generalization
under opponent shift and are not part of \mathcal{D} . The rationale for including O_{medium} and O_{weak} is to

324 introduce opponents of different strengths, which in practice also leads to noticeable variations in
 325 their behavioral patterns. All baselines are trained exclusively on \mathcal{D} , without using additional or
 326 lower-quality data. Further selection criteria are detailed in Appendix G.
 327

328 **Environment.** We evaluate our approach on a diverse set of multi-agent environments from the
 329 PettingZoo suite (Justin K. Terry, 2020b). The benchmark covers four control tasks from the
 330 Multi-Agent Particle Environments (MPE): Tag, Push, Reference, and Spread; two pixel-based
 331 adversarial games from the Atari domain: Tennis and Boxing; and two discrete strategy games
 332 from the Classic category: Connect4 and Texas Hold’em. In addition, we introduce a custom
 333 environment, Badminton (Kuang-Da Wang, 2024b), designed to simulate realistic competitive rallies.
 334 All environments preserve their original observation and reward interfaces, with detailed descriptions
 335 provided in Appendix C.
 336

337 **Evaluation metrics.** Each configuration is evaluated over 1000 episodes with five different random
 338 seeds. For *competitive tasks* (e.g., MPE-Tag, MPE-Push, Atari games, Classic games, and Badminton),
 339 we adopt the win rate as the evaluation metric, defined as $R_{\text{win}} = N_{\text{win}} / (N_{\text{win}} + N_{\text{lose}})$, where N_{win}
 340 and N_{lose} denote the number of wins and losses, respectively.
 341

342 In tasks such as MPE-Tag and MPE-Push, a win is counted when the primary agent’s episodic reward
 343 is greater than or equal to the total reward obtained by all opponents. For *cooperative tasks* (e.g.,
 344 Spread, Reference), we report the average episodic reward of the primary agent, reflecting the overall
 345 team performance under the native reward attribution.
 346

347 5.2 BASELINES

348 In this study, we focus on offline imitation learning and compare it against several state-of-the-art
 349 methods. Offline multi-agent imitation learning is still underexplored, especially diffusion-based
 350 approaches, so we include strong single-agent baselines for comparison, with extensions to multi-
 351 agent settings detailed in Appendix E.2. We select four representative methods:
 352

- 353 • **Behavior Cloning (BC):** Learns a direct state-to-action mapping using supervised learning
 354 without rewards or planning. Each agent is trained independently with a separate BC policy.
 355
- 356 • **Diffusion Behavior Cloning (DBC):** Combines behavior cloning with diffusion-based
 357 generation (Shang-Fu Chen, 2024), minimizing both BC loss and diffusion reconstruction
 358 loss to align actions with expert demonstrations.
 359
- 360 • **Diffusion Policy (DP):** Formulates the policy as a conditional diffusion model, generating
 361 actions by reversing a noise process conditioned on the current state (Tim Pearce, 2023). In
 362 our multi-agent adaptation, the model takes the primary agent’s state as input and outputs
 363 actions for all agents.
 364
- 365 • **Decision Diffusion (DD):** Generates action sequences conditioned on state and re-
 366 turn (Anurag Ajay, 2023). Originally designed for RL with planning, we adapt it to
 367 offline imitation learning without planners or rewards, and also extend it to the multi-agent
 368 setting with the same input–output design as DP.
 369

370 We group baselines by their inference mechanism: diffusion-based methods (DP, DD) and non-
 371 diffusion ones (BC, DBC). Although DBC employs a diffusion module, its action selection follows
 372 standard behavior cloning and thus belongs to the latter group.
 373

374 5.3 STANDARD IMITATION PERFORMANCE RESULTS

375 This experiment evaluates the performance of baseline models trained on the same expert demon-
 376 strations and tested against equally strong opponents O_{expert} . As shown in Table 1, the first row of each
 377 environment (opponent denoted as E) provides the main comparison results.
 378

379 **Comparison with non-diffusion-based policy.** In environments with low-dimensional state spaces
 380 and relatively simple interaction dynamics, such as MPE, BC-based approaches maintain competitive
 381 performance. DBC, which incorporates diffusion-augmented training, achieves notable improvements
 382 on Tag. However, in domains with high-dimensional observations and stochastic transitions, such
 383

378 379 380	Env	Opp.	Non-diffusion-based		Diffusion-based		
			381	382	383	384	385
386	387	388	389	390	391	392	393
Push	E	0.79 ± 0.03	0.79 ± 0.02	0.13 ± 0.03	0.63 ± 0.02	0.81 ± 0.01	
	M	0.77 ± 0.02	0.79 ± 0.03	0.21 ± 0.07	0.61 ± 0.03	0.82 ± 0.02	
	W	0.78 ± 0.07	0.82 ± 0.06	0.15 ± 0.04	0.67 ± 0.03	0.84 ± 0.05	
Tag	E	0.31 ± 0.09	0.38 ± 0.13	0.15 ± 0.03	0.15 ± 0.03	0.33 ± 0.08	
	M	0.42 ± 0.07	0.47 ± 0.05	0.23 ± 0.06	0.43 ± 0.01	0.45 ± 0.03	
	W	0.48 ± 0.02	0.51 ± 0.12	0.30 ± 0.06	0.54 ± 0.10	0.57 ± 0.06	
Spread	E	-11.72 ± 0.34	-11.63 ± 0.36	-15.03 ± 0.41	-13.96 ± 0.25	-11.52 ± 0.41	
	M	-11.93 ± 0.51	-11.86 ± 0.52	-15.41 ± 0.40	-13.83 ± 0.48	-11.87 ± 0.17	
	W	-17.48 ± 0.44	-17.62 ± 0.32	-20.21 ± 0.56	-21.01 ± 0.41	-17.79 ± 0.47	
Reference	E	-30.58 ± 0.82	-30.01 ± 1.16	-28.68 ± 0.86	-27.60 ± 1.02	-26.62 ± 1.07	
	M	-27.87 ± 0.49	-30.13 ± 1.27	-25.65 ± 0.17	-25.01 ± 0.45	-26.81 ± 0.42	
	W	-28.46 ± 0.28	-28.88 ± 0.27	-32.63 ± 0.06	-29.88 ± 0.27	-27.20 ± 0.18	
Connect4	E	0.14 ± 0.02	0.18 ± 0.03	0.12 ± 0.01	0.21 ± 0.02	0.26 ± 0.06	
	M	0.12 ± 0.03	0.22 ± 0.04	0.13 ± 0.05	0.20 ± 0.03	0.41 ± 0.04	
	W	0.41 ± 0.06	0.45 ± 0.03	0.19 ± 0.04	0.29 ± 0.04	0.47 ± 0.07	
Hold'em	E	0.49 ± 0.01	0.53 ± 0.01	0.09 ± 0.02	0.21 ± 0.04	0.55 ± 0.02	
	M	0.53 ± 0.03	0.54 ± 0.04	0.21 ± 0.04	0.35 ± 0.03	0.62 ± 0.02	
	W	0.86 ± 0.05	0.88 ± 0.06	0.34 ± 0.04	0.72 ± 0.03	0.88 ± 0.03	
Tennis	E	0.35 ± 0.07	0.37 ± 0.05	0.62 ± 0.07	0.72 ± 0.06	0.81 ± 0.05	
	M	0.42 ± 0.05	0.46 ± 0.07	0.77 ± 0.05	0.83 ± 0.05	0.90 ± 0.04	
Boxing	E	0.18 ± 0.03	0.17 ± 0.02	0.38 ± 0.02	0.43 ± 0.03	0.47 ± 0.03	
	M	0.15 ± 0.05	0.17 ± 0.08	0.39 ± 0.06	0.45 ± 0.05	0.55 ± 0.04	

400
401 Table 1: The win rate/average reward and standard error across different environments are computed over
402 five seeds. Bold indicates the best result in each row. (E: Expert Opponent, M: Medium Opponent, W: Weak
403 Opponent.)
404

405 as Atari, the performance of BC and DBC degrades considerably. For instance, on Tennis, DDGIL
406 attains a substantially higher win rate compared to BC. The results indicate that DDGIL attains
407 more robust performance in environments with complex observations and interaction dynamics when
408 compared to non-diffusion-based baselines.
409

410 **Comparison with diffusion-based policy.** Compared to DP, which applies a fixed diffusion
411 policy, DDGIL employs a dynamic guidance mechanism that adjusts conditional weighting based
412 on confidence signals from both agent and opponent models. While DP may exhibit stable behavior
413 in certain cases, its static weighting limits adaptability. For instance, DDGIL attains about a 10%
414 higher win rate than DP in Tennis-Expert, highlighting the benefit of adaptive weighting. In contrast,
415 DD relies on state prediction with inverse dynamics, which makes it prone to error accumulation in
416 long-horizon tasks. For example, in Boxing-Expert, DDGIL achieves a relative improvement of about
417 24% over DD, leading to more consistent outcomes. By generating actions directly through denoising
418 and incorporating opponent-aware feedback, DDGIL produces more stable decisions. Notably, in
419 simpler environments such as MPE and Classic games, DDGIL also outperforms both DP and DD,
420 underscoring that the benefit of dynamic guidance extends beyond high-dimensional settings.
421

422 5.4 GENERALIZATION ACROSS OPPONENT STRENGTHS

423 We evaluate each model against alternative opponents: O_{medium} and O_{weak} . This setting tests whether
424 a policy trained on strong opponents can maintain stable performance when faced with unfamiliar or
425 weaker strategies. Results are shown in each environment’s second and third rows in Table 1.
426

427 **Comparison with non-diffusion-based policy.** In MPE tasks, DDGIL performs comparably
428 to DBC and consistently outperforms BC. In cooperative tasks, where weaker opponents act as
429 teammates, its performance shows mixed trends: in Spread-Weak, BC and DBC achieve slightly
430 better scores, but in Reference-Medium, DDGIL clearly surpasses both BC and DBC by 1–2 points.
431 Overall, these results suggest that DDGIL is competitive with non-diffusion baselines and often
432 provides stronger robustness as task difficulty increases.
433

(Agent) vs Opp.	BC	DBC	DD	DP	DDGIL
(K) vs. C	0.14 ± 0.12	0.15 ± 0.11	0.14 ± 0.10	0.43 ± 0.06	0.46 ± 0.12
(K) vs. V	0.22 ± 0.10	0.26 ± 0.11	0.14 ± 0.09	0.61 ± 0.11	0.64 ± 0.12
(C) vs. K	0.07 ± 0.05	0.22 ± 0.12	0.23 ± 0.12	0.39 ± 0.04	0.43 ± 0.17
(C) vs. V	0.14 ± 0.12	0.16 ± 0.12	0.25 ± 0.12	0.70 ± 0.11	0.53 ± 0.13
(V) vs. C	0.08 ± 0.09	0.13 ± 0.11	0.11 ± 0.10	0.35 ± 0.07	0.36 ± 0.12
(V) vs. K	0.06 ± 0.08	0.17 ± 0.10	0.05 ± 0.07	0.36 ± 0.08	0.40 ± 0.12

Table 2: Performance of baseline models in the Badminton environment. The labels K, C, and V correspond to the initials of three real-world players. In the column “(Agent) vs. Opp.”, the player in parentheses denotes the controlled agent, who competes against the other two players in sequence, resulting in six matchups.

Beyond this case, DDGIL consistently outperforms baselines across Atari and Classic environments. For example, in Boxing, BC’s win rate decreases when faced with weaker opponents, DBC shows only a slight improvement, while DDGIL increases by 8% points, demonstrating stronger adaptability. A similar pattern holds in Classic environments: in Connect4, BC decreases by 2% points from Expert to Medium, while DDGIL improves by 15% points.

Comparison with diffusion-based policy. Compared to DP and DD, DDGIL yields higher win-rate improvements against weaker opponents. While diffusion-based models are generally more robust than BC, DP and DD sometimes show non-monotonic reward changes as opponent strength decreases, like Connect4 and Push. In Push, DDGIL achieves the best results. In Boxing, DDGIL maintains consistent performance, whereas DP shows mild fluctuations. These results suggest that fixed guidance (DP) or multi-stage prediction (DD) may not adapt well to interaction shifts.

To understand these results, we analyze why DDGIL achieves stable improvements. At each step, DDGIL combines conditional and unconditional predictions with a confidence-based weight derived from their disagreement (Eqs. 9 and 10). Conditioning remains the default, ensuring the unconditional term never dominates. When disagreement is small, updates stay conservative; when large, the rule amplifies conditional guidance exactly when context matters. In contrast, a fixed weight cannot suit all steps and opponents: high values cause early overreaction, while low values weaken late responses. By recomputing the weight each step, DDGIL adapts guidance strength to observed disagreement, preserving alignment when predictions match and enhancing conditioning when they diverge. This mechanism explains the consistent gains observed across Expert, Medium, and Weak settings.

5.5 EVALUATION IN REAL-WORLD BADMINTON

We extend our evaluation to a virtual sports setting using the Badminton environment. All baseline models are trained on the ShuttleSet dataset (Wei-Yao Wang, 2023a;b), which contains real-world match records. We select three representative players: K, C, and V. The opponent model is RallyNet (Kuang-Da Wang, 2024a), a pretrained imitation agent with diverse playing styles.

Each baseline is evaluated over 20 matches using standard scoring rules. As shown in Table 2, both DP and DDGIL rank among the top performers across all matchups. For instance, in the (K) vs. V scenario, DDGIL achieves a win rate of 0.64, more than triple that of BC and DD. In (C) vs. V, DDGIL yields higher win rates than DBC and DD in most pairings and outperforms DP in 5 of 6 cases, with the only lower result occurring against V.

We make two observations: (1) the hybrid action space, with both discrete and continuous elements, benefits diffusion-based models due to their capacity to model multimodal outputs; and (2) DD underperforms, likely due to long-horizon prediction errors, consistent with its behavior in Atari. **This result also reflects the reward-free adaptation of DD, which removes value gradients that normally guide its long-horizon optimization and can therefore weaken performance under interaction shifts.** While DP remains competitive, DDGIL shows more consistent gains across most matchups. As win rate alone is insufficient, we additionally report interaction traces and trajectory-level statistics (Appendix D.3), including rally length and shot-type transitions.

486

6 CONCLUSION

488 We presented Diffusion Dynamic Guidance Imitation Learning (DDGIL), a diffusion-based frame-
 489 work for offline multi-agent imitation learning. **DDGIL introduces a dynamic guidance rule that**
 490 **adaptively adjusts conditional and unconditional predictions during denoising, enabling stable policy**
 491 **generation under diverse opponent strategies without modifying training.** Experiments across MPE,
 492 Atari, Classic games, and a high-fidelity badminton environment show that DDGIL outperforms
 493 diffusion and non-diffusion baselines, and in badminton it captures tactical patterns of real players,
 494 indicating potential in domains requiring strategic fidelity and adaptability. A key contribution of
 495 DDGIL is the reformulation of diffusion guidance for multi-agent interaction through opponent-aware
 496 conditional scores and an adaptive weighting mechanism, which provides a principled way to adjust
 497 guidance strength in response to varying interaction patterns.

498

BIBLIOGRAPHY

501 Prafulla Dhariwal Alexander Quinn Nichol. Improved denoising diffusion probabilistic models. In *ICML*, 2021.

503 Abhi Gupta Joshua B. Tenenbaum Tommi S. Jaakkola Pulkit Agrawal Anurag Ajay, Yilun Du. Is
 504 conditional generative modeling all you need for decision making? In *ICLR*, 2023.

506 Teng Pang Yan Zhang Yilong Yin Bingzheng Wang, Guoqiang Wu. Diffail: Diffusion adversarial
 507 imitation learning. In *AAAI*, 2024.

508 Yilun Du Zhenjia Xu Eric Cousineau Benjamin Burchfiel Shuran Song Cheng Chi, Siyuan Feng.
 509 Diffusion policy: Visuomotor policy learning via action diffusion. In *RSS*, 2023.

511 Victor R. Lesser Chongjie Zhang. Multi-agent learning with policy prediction. In *AAAI*, 2010.

513 Hyungjin Chung, Jeongsol Kim, Geon Yeong Park, Hyelin Nam, and Jong Chul Ye. CFG++:
 514 manifold-constrained classifier free guidance for diffusion models. In *ICLR 2025*, 2025.

515 Wen Zhang Deheng Ye, Guibin Chen. Towards playing full MOBA games with deep reinforcement
 516 learning. In *NeurIPS*, 2020.

517 Dorsa Sadigh Jiaming Song, Hongyu Ren. Multi-agent generative adversarial imitation learning. In
 518 *NeurIPS*, 2018.

520 Zongqing Lu Jiechuan Jiang. Online tuning for offline decentralized multi-agent reinforcement
 521 learning. In *AAAI*, 2023.

522 Yue Li Jixin Xie, Zhenzhou Shao. Deep reinforcement learning with optimized reward functions for
 523 robotic trajectory planning. *IEEE*, 2019.

525 David Silver Johannes Heinrich. Deep reinforcement learning from self-play in imperfect-information
 526 games. *CoRR*, 2016.

527 Prafulla Dhariwal Alec Radford Oleg Klimov John Schulman, Filip Wolski. Proximal policy
 528 optimization algorithms. *CoRR*, 2017.

530 Pieter Abbeel Jonathan Ho, Ajay Jain. Denoising diffusion probabilistic models. In *NeurIPS*, 2020.

531 Tim Salimans Jonathan Ho. Classifier-free diffusion guidance. 2022.

533 Ananth Hari Justin K. Terry, Benjamin Black. Supersuit: Simple microwrappers for reinforcement
 534 learning environments. *CoRR*, 2020a.

535 Ananth Hari Luis S. Santos Clemens Dieffendahl Niall L. Williams Yashas Lokesh Caroline Horsch
 536 Praveen Ravi Justin K. Terry, Benjamin Black. Pettingzoo: Gym for multi-agent reinforcement
 537 learning. *CoRR*, 2020b.

539 Ping-Chun Hsieh Wen-Chih Peng Kuang-Da Wang, Wei-Yao Wang. Offline imitation of badminton
 player behavior via experiential contexts and brownian motion. *CoRR*, 2024a.

540 Yu-Heng Lin Wei-Yao Wang Wen-Chih Peng Kuang-Da Wang, Yu-Tse Chen. The coachai badminton
 541 environment: Bridging the gap between a reinforcement learning environment and real-world
 542 badminton games. In *AAAI*, 2024b.

543

544 Tengyu Ma Ling Pan, Longbo Huang. Plan better amid conservatism: Offline multi-agent reinforce-
 545 ment learning with actor rectification. In *ICML 2022*, 2022.

546 Shibani Santurkar Dimitris Tsipras Firdaus Janoos Larry Rudolph Aleksander Madry Logan Engstrom,
 547 Andrew Ilyas. Implementation matters in deep policy gradients: A case study on PPO and TRPO.
 548 *CoRR*, 2020.

549

550 Xize Cheng Zhou Zhao Luping Liu, Yi Ren. Diffusion denoising process for perceptron bias in
 551 out-of-distribution detection. *CoRR*, 2022.

552

553 Petru-Daniel Tudosiu Parashkev Nachev Sébastien Ourselin M. Jorge Cardoso Mark S. Graham,
 554 Walter H. L. Pinaya. Denoising diffusion models for out-of-distribution detection. In *CVPR*, 2023.

555

556 Joshua B. Tenenbaum Sergey Levine Michael Janner, Yilun Du. Planning with diffusion for flexible
 557 behavior synthesis. In *ICML*, 2022.

558

559 Umberto Michelucci. An introduction to autoencoders. *CoRR*, 2022.

560

561 Weinan Zhang Yuzheng Zhuang Jun Wang Wulong Liu Yong Yu Minghuan Liu, Ming Zhou. Multi-
 562 agent interactions modeling with correlated policies. In *ICLR, Addis Ababa, Ethiopia, April 26-30,*
 2020, 2020.

563

564 Dhruv Kumar Vivek Myers Erdem Biyik Nathaniel Haynam, Adam Khoja. Multi-agent inverse
 565 q-learning from demonstrations. *CoRR*, 2025.

566

567 Michael Janner Jakub Grudzien Kuba Sergey Levine Philippe Hansen-Estruch, Ilya Kostrikov. IDQL:
 568 implicit q-learning as an actor-critic method with diffusion policies. *CoRR*, 2023.

569

570 Aviv Tamar Jean Harb Pieter Abbeel Igor Mordatch Ryan Lowe, Yi Wu. Multi-agent actor-critic for
 571 mixed cooperative-competitive environments. In *NeurIPS*, pp. 6379–6390, 2017.

572

573 Ming-Hao Hsu Chun-Mao Lai Shao-Hua Sun Shang-Fu Chen, Hsiang-Chun Wang. Diffusion
 574 model-augmented behavioral cloning. In *ICML*, 2024.

575

576 Chang Ye Jeff Braga Dipam Chakraborty Kinal Mehta João G.M. Araújo Shengyi Huang, Rousslan
 577 Fernand Julien Dossa. Cleanrl: High-quality single-file implementations of deep reinforcement
 578 learning algorithms. *Journal of Machine Learning Research*, 2022.

579

580 Linjing Li Jinkyoo Park Shiqi Lei, Kanghoon Lee. Learning strategy representation for imitation
 581 learning in multi-agent games. In *AAAI*, 2025.

582

583 Christian Schröder de Witt Tabish Rashid, Mikayel Samvelyan. QMIX: monotonic value function
 584 factorisation for deep multi-agent reinforcement learning. In *ICML*, 2018.

585

586 Anssi Kanervisto Tim Pearce, Tabish Rashid. Imitating human behaviour with diffusion models. In
 587 *ICLR*, 2023.

588

589 Nicholas Ustaran-Anderegg, Michael Pratt, and Jaime Sabal-Bermudez. AgileRL. URL <https://github.com/AgileRL/AgileRL>.

590

591 David Silver Alex Graves Ioannis Antonoglou Daan Wierstra Martin A. Riedmiller Volodymyr Mnih,
 592 Koray Kavukcuoglu. Playing atari with deep reinforcement learning. *CoRR*, 2013.

593

594 Tsi-Ui Ik Wen-Chih Peng Wei-Yao Wang, Yung-Chang Huang. Shuttleset: A human-annotated
 595 stroke-level singles dataset for badminton tactical analysis. *CoRR*, 2023a.

596

597 Wen-Chih Peng Wei-Yao Wang, Wei-Wei Du. Shuttleset22: Benchmarking stroke forecasting with
 598 stroke-level badminton dataset. *CoRR*, 2023b.

594 Nefeli Andreou Marie-Paule Can Victoria Fernández Abrevaya David Picard Xi Wang, Nicolas Dufour
595 and Vicky Kalogeiton. Analysis of classifier-free guidance weight schedulers. *Trans. Mach. Learn.*
596 *Res.*, 2024.

597 Lantao Yu, Jiaming Song, and Stefano Ermon. Multi-agent adversarial inverse reinforcement learning.
598 In *ICML*, 2019.

599 Xianghao Kong Vagelis Papalexakis Greg Ver Steeg Yunshu Wu, Yingtao Luo. Your diffusion model
600 is secretly a noise classifier and benefits from contrastive training. In *NeurIPS*, 2024.

601 Maryam Zare, Parham M. Kebria, Abbas Khosravi, and Saeid Nahavandi. A survey of imitation
602 learning: Algorithms, recent developments, and challenges. *IEEE Transactions on Cybernetics*,
603 2024.

604 Daochen Zha, Kwei-Herng Lai, Songyi Huang, Yuanpu Cao, Keerthana Reddy, Juan Vargas, Alex
605 Nguyen, Ruzhe Wei, Junyu Guo, and Xia Hu. Rlcard: A platform for reinforcement learning in
606 card games. In *IJCAI*, 2020.

607 Mingyuan Zhou Zhendong Wang, Jonathan J. Hunt. Diffusion policies as an expressive policy class
608 for offline reinforcement learning. In *ICLR*, 2023.

609 Liyuan Mao Bingyi Kang Minkai Xu Yong Yu Stefano Ermon Weinan Zhang Zhengbang Zhu,
610 Minghuan Liu. Madiff: Offline multi-agent learning with diffusion models. In *NeurIPS*, 2024.

611 Jianye Hao Fei Ni Yi Ma Pengyi Li Yan Zheng Zibin Dong, Yifu Yuan. Cleandiffuser: An easy-to-use
612 modularized library for diffusion models in decision making. *arXiv preprint arXiv:2406.09509*,
613 2024. URL <https://arxiv.org/abs/2406.09509>.

614

615

616

617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

643

644

645

646

647

648
649
650
651
652

Appendix

653	A Theoretical Insights and Comparative Analysis	13
654	A.1 Verification of the Dynamic Weighting	14
655	A.2 Relationship Between Diffusion Policy (DP) and DDGIL	15
656		
657		
658	B Algorithm	16
659	B.1 Embedding Model Algorithm	16
660	B.2 Diffusion Dynamic Guidance Policy Algorithm	16
661		
662		
663	C Environment settings	17
664		
665	D Extra Experiments	19
666	D.1 Effect of varying dataset size	19
667	D.2 Fixed guidance weights versus adaptive weights	20
668	D.3 Analysis of performance in Badminton	21
669	D.4 Compare with Reinforcement Learning models	22
670	D.5 Scalability of multiple agents for DDGIL	22
671	D.6 Necessity of multiple diffusion models in DDGIL	23
672		
673		
674		
675	E Model architecture	24
676	E.1 Expert Reinforcement Learning Model	24
677	E.2 Baseline model	25
678	E.3 DDGIL model	27
679	E.4 Computing Resources	28
680		
681		
682		
683	F Inference Details	29
684		
685	G Dataset information	29
686	G.1 Dataset Construction	29
687	G.2 Dataset Collection in Different Types of Environments	30
688		
689		
690	H The Use of Large Language Models (LLMs)	31
691		
692	I Limitations and Discussion	31
693		
694		
695		
696		
697	A THEORETICAL INSIGHTS AND COMPARATIVE ANALYSIS	
698		
699	This appendix first examines the rationale of the dynamic weight computation and its mathematical	
700	consistency. We then illustrate how the mechanism balances stability across different conditions.	
701	Finally, we contrast the resulting formulation with Diffusion Policy (DP) to highlight the methodological	
	differences between DP and DDGIL.	

702 A.1 VERIFICATION OF THE DYNAMIC WEIGHTING
703

704 To motivate the proposed guidance mechanism, it is necessary to clarify why $d_t = \|\hat{\epsilon}_t^B - \hat{\epsilon}_t^G\|$ with
705 $w_D^{(t)} = \sigma(d_t)$ constitutes a reasonable choice of dynamic weight. This formulation directly links the
706 level of disagreement between predictors to the strength of conditional guidance, ensuring that the
707 weighting is adapted in a principled and mathematically consistent manner.

708
709 **Setup and notations.** Fix a reverse diffusion step t and state x_t . Let $\hat{\epsilon}_t^G$ denote the conditional noise
710 prediction (given the current condition), $\hat{\epsilon}_t^B$ the opponent prediction (e.g., averaged over opponents),
711 and $\hat{\epsilon}_t^{NG}$ the unconditional prediction. Note that the main text used the notations $\hat{\epsilon}_G$, $\hat{\epsilon}_G^N$, and $\hat{\epsilon}_B$;
712 here we include the timestep index t explicitly for clarity of exposition. Define

$$713 \quad d_t := \|\hat{\epsilon}_t^B - \hat{\epsilon}_t^G\|_2, \quad w_D^{(t)} := \sigma(d_t) \in [0.5, 1],$$

714 where $\sigma(u) = \frac{1}{1+e^{-u}}$ is the logistic sigmoid. The following analysis establishes $d_t \mapsto w_D^{(t)}$ as a
715 principled and stable dynamic weighting rule.

716
717 **(i) From noise prediction to score.** Introduce the score notation

$$718 \quad s_t^G(x_t) := \nabla_{x_t} \log p_G(x_t), \quad s_t^B(x_t) := \nabla_{x_t} \log p_B(x_t).$$

719 Under the DDPM parameterization, there exists a constant $c_t > 0$ (dependent only on the noise
720 schedule) such that

$$721 \quad s_t^G(x_t) \approx -c_t \hat{\epsilon}_t^G, \quad s_t^B(x_t) \approx -c_t \hat{\epsilon}_t^B. \quad (11)$$

722
723 **(ii) d_t as the magnitude of a log-likelihood ratio gradient.** Define the stepwise log-likelihood
724 ratio

$$725 \quad \ell_t(x_t) := \log \frac{p_G(x_t)}{p_B(x_t)}.$$

726 Differentiating with respect to x_t and applying Eq. 11 gives

$$727 \quad \nabla_{x_t} \ell_t(x_t) = s_t^G(x_t) - s_t^B(x_t) \\ 728 \quad \approx -c_t (\hat{\epsilon}_t^G - \hat{\epsilon}_t^B), \quad (12)$$

729 and thus

$$730 \quad \|\nabla_{x_t} \ell_t(x_t)\| \approx c_t \|\hat{\epsilon}_t^G - \hat{\epsilon}_t^B\| = c_t d_t. \quad (13)$$

731 Hence d_t (up to scale) coincides with the gradient magnitude of a discriminative objective ℓ_t : large
732 values indicate divergent explanations from p_G and p_B , while small values indicate agreement.

733
734 **(iii) Smooth mapping via $\sigma(\cdot)$.** Since $d_t \geq 0$, the logistic map yields

$$735 \quad w_D^{(t)} = \sigma(d_t) \in [0.5, 1],$$

$$736 \quad \frac{\partial w_D^{(t)}}{\partial d_t} = \sigma(d_t)(1 - \sigma(d_t)) > 0, \quad (14)$$

$$737 \quad \frac{\partial^2 w_D^{(t)}}{\partial d_t^2} = \sigma(d_t)(1 - \sigma(d_t))(1 - 2\sigma(d_t)).$$

738 Thus $w_D^{(t)}$ increases monotonically with d_t , remains bounded within $[0.5, 1]$, and varies smoothly
739 across steps.

740
741 **(iv) One-step consistency via line integral of $\nabla \ell_t$.** The deterministic mean displacement is

$$742 \quad \Delta x_t := x_{t-1} - x_t = -\kappa_t \hat{\epsilon}_t, \quad \kappa_t := \frac{1 - \alpha_t}{\sqrt{\alpha_t} \sqrt{1 - \bar{\alpha}_t}}, \quad (15)$$

743 with dynamically mixed estimate

$$744 \quad \hat{\epsilon}_t = w_D^{(t)} \hat{\epsilon}_t^G + (1 - w_D^{(t)}) \hat{\epsilon}_t^{NG}. \quad (16)$$

756 Along the path $\gamma(\tau) = x_t + \tau \Delta x_t$, $\tau \in [0, 1]$, the change in ℓ_t is
 757

$$\begin{aligned} 758 \quad \Delta \ell_t &= \int_0^1 \nabla_{x_t} \ell_t(\gamma(\tau))^{\top} \gamma'(\tau) d\tau \approx \nabla_{x_t} \ell_t(x_t)^{\top} \Delta x_t \\ 759 \\ 760 &\stackrel{(\text{Eq.12, 15})}{\approx} \kappa_t c_t (\hat{\epsilon}_t^G - \hat{\epsilon}_t^B)^{\top} [w_D^{(t)} \hat{\epsilon}_t^G + (1 - w_D^{(t)}) \hat{\epsilon}_t^{NG}]. \\ 761 \end{aligned} \quad (17)$$

762 Expanding gives
 763

$$\Delta \ell_t = \kappa_t c_t \left\{ w_D^{(t)} (\hat{\epsilon}_t^G - \hat{\epsilon}_t^B)^{\top} \hat{\epsilon}_t^G + (1 - w_D^{(t)}) (\hat{\epsilon}_t^G - \hat{\epsilon}_t^B)^{\top} \hat{\epsilon}_t^{NG} \right\}. \quad (18)$$

766 When the conditional direction aligns better with $\nabla \ell_t$ than the unconditional one, a larger $w_D^{(t)}$ yields
 767 a greater increase in ℓ_t . Because $w_D^{(t)}$ is monotone in d_t (Eq. 14) and $d_t \propto \|\nabla \ell_t\|$ (Eq. 13), the
 768 weighting adapts emphasis toward the conditional signal when discriminative evidence is strong, and
 769 away when it is weak.

770 **(v) Stability and compatibility.** Since $w_D^{(t)} \in [0.5, 1]$, the update in Eq. 16 remains a convex
 771 combination with conditional dominance. The derivative $\sigma'(d_t)$ in Eq. 14 bounds sensitivity to d_t ,
 772 reducing variability due to noisy estimates. Together with Eq. 17, this shows that the rule increases ℓ_t
 773 proportionally to evidence while preserving the standard DDPM update form.

774 **Summary.** The analysis highlights four properties. First, $d_t = \|\hat{\epsilon}_t^B - \hat{\epsilon}_t^G\|$ corresponds (up
 775 to c_t) to the gradient magnitude of ℓ_t (Eq. 13), linking it to local model disagreement. Second,
 776 $w_D^{(t)} = \sigma(d_t)$ is monotone and smooth (Eq. 14), assigning higher weight under stronger evidence.
 777 Third, bounding $w_D^{(t)}$ within $[0.5, 1]$ enforces convex mixing and moderates sensitivity. This design
 778 keeps the conditional branch relatively emphasized over the unconditional component; however,
 779 while the text describes small d_t as being close to the unconditional update, the lower bound of 0.5
 780 implies that the update only partially approaches the unconditional case. Finally, the line-integral
 781 argument (Eqs. 15–18) shows compatibility with the reverse denoising update. These properties
 782 together support d_t and $w_D^{(t)}$ as a principled dynamic weighting rule.

785 A.2 RELATIONSHIP BETWEEN DIFFUSION POLICY (DP) AND DDGIL

786 This section highlights the differences between our method DDGIL and Diffusion Policy (DP).
 787 Among related works, DP is the closest to ours, as it also employs a diffusion model for behavior
 788 cloning. If the dynamic guidance mechanism is removed, the opponent model is discarded, and the
 789 guidance weight w is fixed to a constant (e.g., $w_D = 1$), our method reduces to DP.

790 In DP, a single conditional diffusion model is trained to predict actions from observations or auxiliary
 791 information (e.g., state trajectories or history). By contrast, DDGIL introduces an auxiliary model
 792 to capture opponent behaviors and dynamically rebalances the contributions of the primary agent
 793 and opponent predictions at each denoising step, conditioned on the agent’s observation. The final
 794 predicted noise is given by:
 795

$$\text{DP: } \hat{\epsilon} = w \cdot \hat{\epsilon}_G + (1 - w) \cdot \hat{\epsilon}_G^N, \quad w \in [0, 1] \text{ (constant)} \quad (19)$$

$$\text{DDGIL: } \hat{\epsilon} = w_D \cdot \hat{\epsilon}_G + (1 - w_D) \cdot \hat{\epsilon}_G^N, \quad w_D = \sigma(\|\hat{\epsilon}_B - \hat{\epsilon}_G\|_2) \quad (20)$$

800 **Key differences.** (i) DP relies on a fixed weight w in Eq. 19, whereas DDGIL replaces it with
 801 a dynamic w_D in Eq. 20. (ii) DDGIL increases the conditional component when the discrepancy
 802 between $\hat{\epsilon}_B$ and $\hat{\epsilon}_G$ grows ($d \uparrow \Rightarrow w_D \uparrow$), thereby prioritizing state-conditioned cues over the
 803 unconditional prior. This mitigates collapse to recurring patterns (e.g., repeated mid-court clears in
 804 Badminton) and improves robustness under opponent or style shifts. (iii) DDGIL operates purely at
 805 inference, requiring neither retraining nor architectural modification. (iv) **Unlike opponent-policy
 806 models in strategic reasoning frameworks, our opponent diffusion models do not generate actions
 807 and are used only to provide contextual score predictions for guiding the primary agent.**

808 Moreover, DP implicitly assumes that the conditional distribution $p(a | s)$ is sufficient for imitation,
 809 independent of interactive context. Our findings suggest that this assumption is inadequate in multi-
 810 agent environments, where strategic variability is strongly influenced by other agents. By explicitly

810 modeling the opponent and incorporating its predictions as an adaptive reference, DDGIL improves
 811 both robustness and behavioral fidelity. Main experiments (see Section 5.3) confirm that this dynamic
 812 formulation consistently outperforms DP, underscoring the importance of opponent-aware guidance
 813 and adaptive inference.
 814
 815

816 B ALGORITHM

817 B.1 EMBEDDING MODEL ALGORITHM

818 Our overall training pipeline focuses primarily on training a diffusion model for each agent. However,
 819 in specific environments, the observed state is not a low-dimensional vector but a high-dimensional
 820 image. This differs from array-based environments such as MPE, Texas Hold’em, and Badminton,
 821 where the state can be directly fed into the model. For example, Atari environments provide
 822 $84 \times 84 \times 6$ compressed image frames, and Connect4 uses a $7 \times 6 \times 2$ binary tensor representation.
 823 Since diffusion models expect vectorized conditional inputs, using such image-based states can lead
 824 to training instability or poor convergence.
 825

826 To address this, we train an auxiliary embedding model f_{emb} to transform high-dimensional states
 827 into vector representations before they are used as diffusion conditions (Michelucci, 2022). This
 828 embedding model is pretrained once and kept fixed during all downstream tasks. The overall
 829 procedure is summarized in Algorithm 1. To ensure a fair comparison, we apply the same pre-trained
 830 f_{emb} across all baselines, including BC, DP, DBC, and DD, in environments that involve image-based
 831 states (Atari and Connect4). The architecture and parameter settings of f_{emb} , as well as all other
 832 models, are detailed in Appendix E.
 833
 834

835 B.2 DIFFUSION DYNAMIC GUIDANCE POLICY ALGORITHM

836 We present the dynamic guidance mechanism used during inference. For environments requiring state
 837 embeddings, a pretrained embedding model encodes state representations for training and inference.
 838 Additionally, the dynamic guidance mechanism relies on pretrained diffusion models for both the
 839 primary agent and the opponent agents. The algorithm is summarized in Algorithm 2.
 840

841 Algorithm 1 Training of Agent-Specific Embedding Model f_{emb}

```

842 1: Input: Dataset  $\mathcal{D} = \{s^i\}$ , encoder  $E^i$ , decoder  $D^i$  for each agent  $i$ 
843 2: Output: Trained encoder-decoder pairs  $(E^i, D^i)$ 
844 3: for each agent  $i \in \{1, \dots, K\}$  do
845 4:   Initialize encoder  $E^i$  and decoder  $D^i$ 
846 5: end for
847 6: while not converged do
848 7:   Sample batch of raw image states or one-hot array  $\{s_t^i\}$  from  $\mathcal{D}$  for all agents
849 8:   for each agent  $i$  do
850 9:     if  $s_i$  is an image then:
851 10:      Normalize input:  $\tilde{s}^i \leftarrow s^i / 255.0$ 
852 11:      else
853 12:        input:  $\tilde{s}^i \leftarrow s^i$ 
854 13:      end if
855 14:      Encode:  $z^i \leftarrow E^i(\tilde{s}^i)$ 
856 15:      Decode:  $\hat{s}^i \leftarrow D^i(z^i)$ 
857 16:      Compute reconstruction loss:  $\mathcal{L}_i \leftarrow \|\hat{s}^i - \tilde{s}^i\|^2$ 
858 17:      Update  $E^i$  and  $D^i$  using gradient of  $\mathcal{L}_i$ 
859 18: end for
860 19: end while

```

864 **Algorithm 2** Inference with Dynamic Guidance

```

866 1: Input: Denoising models  $\epsilon_\theta^G, \epsilon_\phi^{B_i}$ , Diffusion steps  $T$ , Embedding model  $f_{\text{emb}}$ , Primary Agent
867   index  $G$ , Opponent index  $B$ , Number of opponents  $k$ , Primary Agent's state  $s^G$ 
868 2: Output: Primary agent's action  $a^G$ 
869 3: if  $s^G$  is an image then
870   4:   Embedding  $z \leftarrow f_{\text{emb}}(s^G)$ 
871 5: else
872   6:    $z \leftarrow s^G$ 
873 7: end if
874 8: Initialize  $x_T \sim \mathcal{N}(0, I)$ 
875 9: for  $t = T, \dots, 1$  do
876   10: Predict noise with condition  $\hat{\epsilon}_G \leftarrow \epsilon_\theta^G(x_t, z, t)$ 
877   11: Predict noise without condition  $\hat{\epsilon}_G^N \leftarrow \epsilon_\theta^G(x_t, \emptyset, t)$ 
878   12: for each  $i \in \{1, \dots, k\}$  do
879     13:   Predict noise with condition  $\hat{\epsilon}_{B_i} \leftarrow \epsilon_\phi^{B_i}(x_t, z, t)$ 
880   14: end for
881   15: Compute the mean of the conditional noise  $\hat{\epsilon}_B = \mathbb{E}[\hat{\epsilon}_{B_i}]$ 
882   16: Compute confidence  $d \leftarrow \|\hat{\epsilon}_B - \hat{\epsilon}_G\|$ 
883   17: Compute dynamic guidance  $w_D \leftarrow \sigma(d)$   $\triangleright \sigma(x) = \frac{1}{1+e^{-x}}$ 
884   18: Compute epsilon  $\hat{\epsilon} \leftarrow w_D \cdot \hat{\epsilon}_G + (1 - w_D) \cdot \hat{\epsilon}_G^N$ 
885   19:  $(\mu_t, \Sigma_t) \leftarrow \text{Denoise}(x_t, \hat{\epsilon})$ 
886   20:  $x_{t-1} \sim \mathcal{N}(\mu_t, \Sigma_t)$ 
887 21: end for
888 22: Action selection: For discrete dimensions,  $a^G \leftarrow \arg \max_a (x_0[a])$ ; for continuous dimensions,
889   take the real-valued output from  $x_0$ .

```

890 C ENVIRONMENT SETTINGS

892 To comprehensively evaluate the adaptability and robustness of our proposed method under diverse
893 interaction scenarios, we select a range of representative multi-agent reinforcement learning (MARL)
894 environments as our evaluation benchmarks. A key criterion for environment selection is a multi-
895 agent structure; accordingly, we primarily adopt environments from the PettingZoo library (Justin
896 K. Terry, 2020b). PettingZoo is a Python library designed explicitly for MARL research, offering
897 a unified API and supporting various interaction types, including cooperative, competitive, and
898 communication-based settings.

900 Environment	901 Category	902 Interaction Type	903 Action Type
902 Push	903 MPE	904 Parallel	905 Discrete
903 Tag	904 MPE	905 Parallel	906 Discrete
904 Spread	905 MPE	906 Parallel	907 Discrete
905 Reference	906 MPE	907 Parallel	908 Discrete
906 Tennis	907 Atari	908 Parallel	909 Discrete
907 Boxing	908 Atari	909 Parallel	910 Discrete
908 Connect4	909 Classic	910 AEC	911 Discrete
909 Texas Hold'em	910 Classic	911 AEC	912 Discrete
910 Badminton	911 Real	912 AEC	913 Discrete + Continuous

912 Table 3: Metadata for each environment, including its category, interaction mode, and action space type.
913 Badminton includes hybrid action/state features (state: 11 discrete + 6 continuous; action: 11 discrete + 4
914 continuous).

915 PettingZoo environments are categorized into two major interaction protocols: Agent Environment
916 Cycle (AEC) and Parallel. AEC environments enforce turn-based interactions where agents observe
917 and act sequentially, making them suitable for step-wise strategic settings. In contrast, Parallel

environments allow all agents to observe and act, simultaneously simulating real-time or synchronous interaction. These two formats differ in data collection structure, and we detail their respective recording formats in the Dataset section.

To capture diverse task structures and input modalities, we categorize the environments into four types: (1) **MPE** (Multi-agent Particle Environments), which are vector-based and emphasize coordination and adversarial interaction; (2) **Atari**, which provides pixel based inputs and complex competitive dynamics; (3) **Classic** environments, such as Connect4 and Texas Hold'em, which feature well-defined rules and game-theoretic structure; and (4) **Badminton**, a high-fidelity sports simulation inspired by real-world gameplay.

Table 3 summarizes the structural and behavioral characteristics of each domain, including category, interaction type, and action space. Additional specifications, such as state dimensionality, number of agents, action space size, and maximum episode length, are provided in Table 4, which also guides model configuration and training.

- **MPE-Tag.** This predator-prey environment involves one fast-moving good agent and three slower adversaries. The good agent incurs a penalty of -10 upon each collision with an opponent, while adversaries receive a reward of +10 for successfully hitting the good agent. The environment also includes two static obstacles blocking movement and influencing path planning. The good agents don't run to infinity, and they are also penalized for exiting the area
- **MPE-Push.** This environment consists of one good agent, one opponent, and a single landmark. The good agent receives a reward based on its proximity to the landmark, while the opponent is rewarded proportionally to the difference between its distance and the good agent's distance to the landmark. As a result, the opponent is incentivized to push the good agent away from the landmark to maximize its reward.
- **MPE-Spread.** This environment consists of N agents and N landmarks (with a default of $N = 3$). Agents are tasked with collectively covering all landmarks while minimizing inter-agent collisions. Globally, the team receives a shared reward based on the sum of the minimum distances from each landmark to the nearest agent. Locally, each agent incurs a penalty of -1 for every collision with another agent. The relative contribution of global versus local rewards is modulated by a local ratio, allowing for flexible trade-offs between cooperation and collision avoidance.
- **MPE-Reference.** This environment features two agents and three uniquely colored landmarks. Each agent aims to reach its designated target landmark, the identity of which is known only to the other agent. Both act as speakers and listeners, exchanging information to locate their targets. Local rewards are based on each agent's distance to its target, while global rewards depend on the average distance of all agents to their respective targets.
- **Atari-Tennis.** This environment is a competitive two-player game focused on positioning and prediction. Each agent aims to strike the ball past the opponent while preventing it from crossing their own side. A successful point yields a reward of +1 to the scorer and -1 to the opponent. To prevent stalling, players are penalized -1 if they fail to serve within 3 seconds of receiving the ball, introducing a non-zero-sum aspect to the game.
- **Atari-Boxing.** This environment simulates an adversarial boxing match emphasizing precise control and reactive strategy. Over a fixed duration of approximately 128 steps, agents can move and punch at each timestep. Scoring is based on the effectiveness of punches: 1 point for a long jab, 2 points for a close-range power punch. Each successful action yields a corresponding positive reward for the scorer and an equivalent negative reward for the opponent.
- **Classic-Connect4.** It is a two-player, turn-based game that aims to align four consecutive tokens vertically, horizontally, or diagonally on a 7-column grid. On each turn, a player drops a token into a selected column, and it falls to the lowest available position. Tokens cannot be placed in full columns. The game ends when either player achieves a sequence of four tokens or when all columns are filled, resulting in a draw if no player has won.
- **Classic-Texas Hold'em (Limit).** It is a simplified variant of Limit Texas Hold'em with two players, two betting rounds, and a deck of six cards (Jack, Queen, King in two suits). Each

972
 973
 974
 975
 976
 977
 978
 979
 980
 981
 982
 983
 984
 985
 986
 987
 988
 989
 990
 991
 992
 993
 994
 995
 996
 997
 998
 999
 1000
 1001
 1002
 1003
 1004
 1005
 1006
 1007
 1008
 1009
 1010
 1011
 1012
 1013
 1014
 1015
 1016
 1017
 1018
 1019
 1020
 1021
 1022
 1023
 1024
 1025

player is dealt one private card, followed by a betting round, after which a single public card is revealed. A second round of betting ensues. The player with the highest-ranked hand at the end wins the game and receives a reward of $+1$, while the loser gets -1 . At any point, a player may fold, forfeiting the game.

- **Real-World Sports: Badminton.** The Badminton environment is sourced from the CoachAI framework (Kuang-Da Wang, 2024b) and built upon the ShuttleSet dataset (Wei-Yao Wang, 2023b). ShuttleSet is currently the largest publicly available badminton singles dataset, covering 44 matches from 2018 to 2021 with over 36,000 annotated strokes. Each rally is annotated with rich metadata, including shot type, shot location, and player positions, making it especially suitable for imitation learning and tactical modeling. Each agent controls a player in a fast-paced rally, aiming to return shots and score against the opponent. This environment introduces real-world game dynamics, long-term interaction, and a more complex state-action distribution than synthetic settings.

Environment	Agent SD	Opp. SD	Num Agent	Num Opp.	AD	Traj. Length
Push	19	8	1	1	5	25
Tag	14	16	1	3	5	25
Spread	18	–	3	0	5	25
Reference	21	–	2	0	50	25
Connect4	$2 \times 6 \times 7$	$2 \times 6 \times 7$	1	1	7	42 (21)
Hold'em	72	72	1	1	4	50 (25)
Boxing	$84 \times 84 \times 6$	$84 \times 84 \times 6$	1	1	18	128
Tennis	$84 \times 84 \times 6$	$84 \times 84 \times 6$	1	1	18	128
Badminton	17	17	1	1	15	60 (30)

Table 4: Environment-specific state/action dimensions and agent settings. “Traj. Length” denotes the maximum steps per episode. “SD” and “AD” denote the abbreviations for State Dimension and Action Dimension, respectively. Values in parentheses indicate per-agent lengths in turn-based (AEC) settings, typically computed as max steps divided by the number of agents. This table supports the model architecture in Table 3.

D EXTRA EXPERIMENTS

This section presents an ablation study to investigate the performance differences between DDGIL and other baselines. Specifically, we examine: (i) the impact of varying the quantity of training data; (ii) the effect of using fixed weights compared to DDGIL; (iii) the adequacy of win rate in the Badminton environment, supplemented by analyses of score causes, mistake patterns, and rally length distributions to assess imitation quality; (iv) comparisons with reinforcement learning models; (v) the scalability of DDGIL to multiple agents; and (vi) the necessity of employing multiple diffusion models within DDGIL.

D.1 EFFECT OF VARYING DATASET SIZE

We begin by analyzing the effect of data quantity on model performance. The dataset sizes in the main experiments were set to 500 episodes for MPE and Classic, and 250 for Atari. We exclude the Badminton environment from this analysis, as its dataset is fixed and does not allow for controlled variation in data volume. To evaluate sensitivity to data availability, we train each baseline with varying amounts of data: 50, 100, and 250 episodes for MPE and Classic; 10, 50, and 100 episodes for Atari. All models are trained using the same configurations as in the main experiments and evaluated against the opponent at the expert level O_{expert} . The results are presented in Figure 3.

The figure shows that most baseline models improve as the training data increases. BC, DBC, and DDGIL achieve competitive results in the three MPE tasks, whereas DDGIL consistently outperforms other baselines in Atari and Classic environments. Moreover, the average reward improvement across all tasks remains within 50% when comparing the smallest and largest dataset sizes. These results suggest that, despite some variations, the overall performance of our models is relatively robust with respect to the quantity of data.

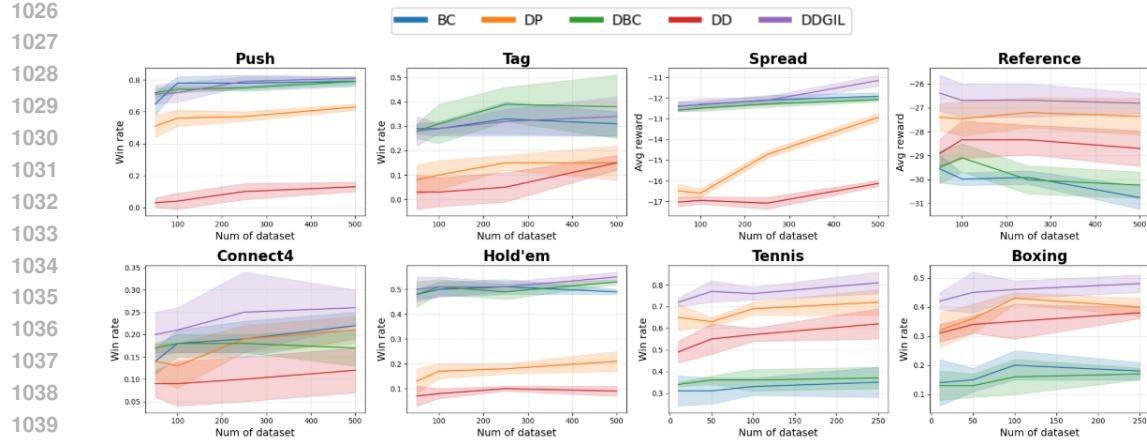


Figure 3: Dataset configurations for analyzing the impact of data quantity on model performance. MPE and Classic tasks are evaluated with 50, 100, and 250 episodes, while Atari tasks are evaluated with 10, 50, and 100 episodes.

D.2 FIXED GUIDANCE WEIGHTS VERSUS ADAPTIVE WEIGHTS

We evaluate the effect of fixed guidance weights by setting w to 1.0, 0.75, 0.5, 0.25, and compare these fixed-weight settings against the adaptive weighting used in DDGIL. Here, $w = 1.0$ corresponds to agent-only guidance (equivalent to DP), while $w = 0.0$ uses only the opponent policy. All models are evaluated with three random seeds over 300 episodes. Results are shown in Figure 4.

In MPE environments, fixing w above 0.75 achieves higher rewards than DDGIL, with Push and Tag showing improvements of 0.2 to 0.3. For other tasks, $w = 1.0$ performs close to DDGIL but remains slightly lower. In Hold'em, the best performance occurs at $w = 0.5$, suggesting a balance between agent and opponent guidance benefits policy learning. Unlike Push or Tag, Hold'em relies on adapting to opponent strategies, making mixed guidance more effective than purely agent-driven decisions. When $w = 0.0$, performance consistently degrades across all tasks. Despite sharing the same noise and denoising process, the opponent policy optimizes toward opponent behavior, which conflicts with reproducing the agent's strategy. This results in a reward gap that is roughly two to three times larger than in other settings.

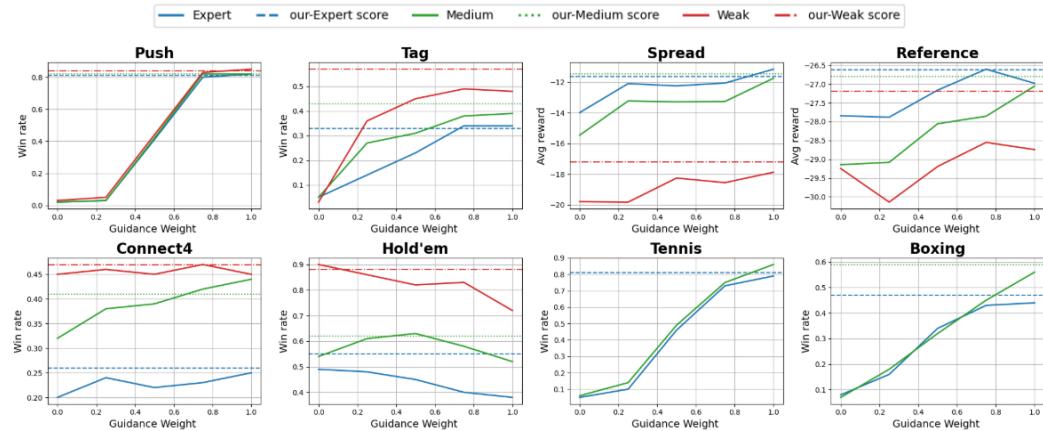


Figure 4: Performance comparison between fixed guidance weights w and adaptive DDGIL. Each setting is evaluated over 300 episodes, with $w = 1.0$ representing agent-only guidance and $w = 0.0$ representing opponent-only guidance.

1080
1081

D.3 ANALYSIS OF PERFORMANCE IN BADMINTON

1082
1083
1084
1085
1086
1087

In Section 5.5, we compared baseline models based on win rates. However, given badminton’s interactive and dynamic nature, win rate alone is insufficient to evaluate how well a model replicates player behavior. To address this, we further analyze the quality of generated match processes by examining rally length distribution and score-related landing positions. The analysis uses data from 20 matches recorded during interactions with real players, following standard badminton rules, including match point settings.

1088
1089
1090
1091

Rally Length Distribution. Rally length reflects the tempo of exchanges and error control, serving as a key indicator of realistic gameplay. Models that fail to capture proper shot selection and defensive reactions often produce abnormally short rallies, dominated by serve or receive errors.

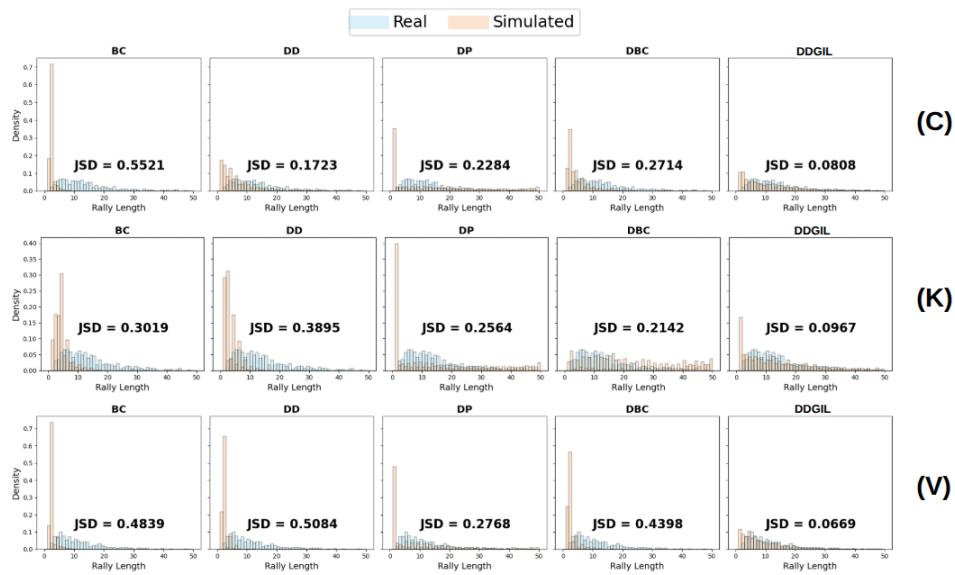
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
11101111
1112
1113
1114

Figure 5: Comparison of rally length distributions between real players and generated simulations. Each subfigure shows the distribution for a specific player (C, K, V) across different models. JSD quantifies the difference between real and simulated distributions. DDGIL achieves the lowest JSD in all cases, indicating superior replication of realistic rally patterns.

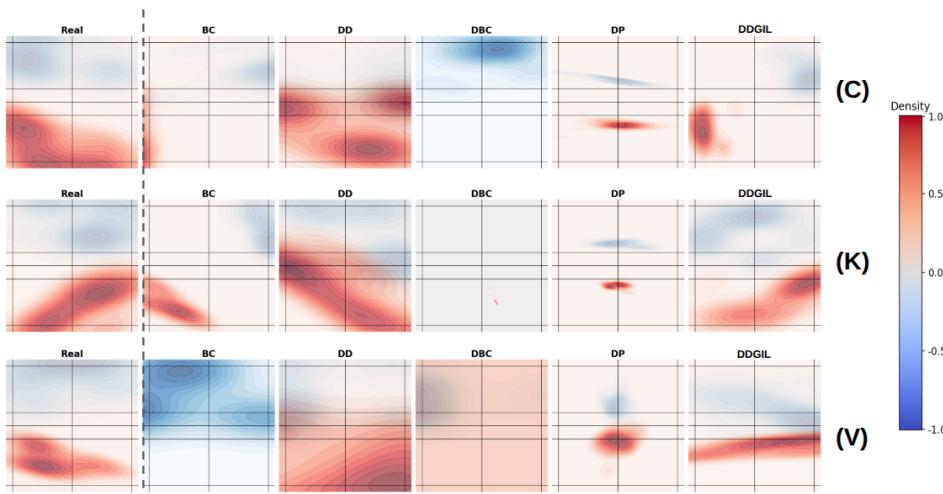
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

Figure 6: Landing position distributions of scoring (blue) and losing (red) shots for each player (C, K, V) across different models. Distributions are visualized using KDE heatmaps.

1134 To evaluate this, we compute the rally length distribution for each model and measure its divergence
 1135 from real data using Jensen–Shannon divergence (JSD), as shown in Figure 5. DDGIL achieves
 1136 the smallest divergence, with an average JSD below 0.1, while other baselines exceed 0.2. Most
 1137 alternative models show rallies concentrated between 1 to 3 shots, typically due to serve faults (length
 1138 1) or failed returns (length 2). Further inspection reveals these errors stem from incorrect shot types
 1139 or boundary violations, highlighting DDGIL’s superior ability to replicate realistic shot selection and
 1140 play styles.

1141

1142 **Win/Lose landing Distribution.** We also analyze the landing positions of scoring and losing shots.
 1143 Specifically, we record the positions where a point is won (on the opponent’s court) and where a
 1144 point is lost (on the agent’s court), excluding out-of-bounds shots. This assesses whether the model
 1145 captures the player’s preferred attack zones and common defensive weaknesses. Alignment in scoring
 1146 positions indicates learned offensive tendencies, while consistency in losing positions reflects an
 1147 understanding of a player’s typical weaknesses.

1148 Figure 6 presents the comparison, with red areas denoting losing shot distributions and the blue
 1149 regions for scoring shots. Though no numerical metrics are provided, KDE visualizations show
 1150 that DDGIL’s distributions closely match those of the real players, particularly for players K and
 1151 C. For player V, the alignment is less precise but still reasonable. In contrast, BC and DBC exhibit
 1152 imbalanced patterns, often lacking losing shot distributions due to frequent serve faults and out-of-
 1153 bounds hits, which are not reflected in the landing statistics. DP and DD show distinct behaviors: DD
 1154 produces overly dispersed landing positions, lacking clear attack patterns, while DP focuses on fixed
 1155 regions for both scoring and losing shots, indicating limited adaptability. These observations further
 1156 demonstrate DDGIL’s advantage in mimicking player-specific strategies and behaviors.

1157

1158 D.4 COMPARE WITH REINFORCEMENT LEARNING MODELS

1159

1160 In this work, although our primary setting is reward-free, we additionally compare DDGIL with offline
 1161 RL methods to assess its generalization and stability. For a fair comparison, we recollected datasets
 1162 with reward information and included IDQL (Philippe Hansen-Estruch, 2023) and OMAR (Ling Pan,
 1163 2022) as baselines. All evaluations follow the same protocol as in the main experiments under Expert,
 1164 Medium, and Weak opponent conditions.

1165

1166 Table 5 reports the comparison between DDGIL and offline RL baselines. On MPE tasks, DDGIL
 1167 achieves consistent gains, with win rates on Push and Tag exceeding IDQL and OMAR by up to
 1168 +6–8% under certain opponent conditions. In Spread, DDGIL remains competitive but trails the
 1169 strongest baseline by roughly 2–3%. In Classic and Atari domains, RL methods generally dominate:
 1170 in Connect4 and Tennis, DDGIL is lower by 5–10%, while in Boxing the gap narrows to within
 2–3%. For Hold’em, DDGIL surpasses RL baselines under Weak.

1171

1172 Overall, these results indicate that DDGIL is particularly effective in MPE environments, where
 1173 coordination and adaptation to opponent strategies are critical. While RL baselines retain an advantage
 1174 in Connect4 and most Atari domains, DDGIL demonstrates competitive performance in reward-free
 1175 settings and achieves superior outcomes in interaction-heavy tasks such as Push, Tag, and Expert
 Hold’em.

1176

1177 D.5 SCALABILITY OF MULTIPLE AGENTS FOR DDGIL

1178

1179 To evaluate the scalability of DDGIL with respect to the number of opponents, we conducted
 1180 experiments on the MPE Tag task with 3, 6, and 10 opponents. Each configuration was evaluated
 1181 in a fixed expert setting for 100 episodes, averaging the results in three random seeds with standard
 1182 deviations, as reported in Table 6. In addition, single-step inference time and memory usage were
 1183 recorded.

1184

1185 Results show that DDGIL maintains stable win rates as the number of opponents increases. With
 1186 three opponents, DDGIL achieves a win rate of 0.34, which is lower than DBC. As the number of
 1187 opponents increases to six and ten, DDGIL maintains the highest performance, whereas BC and DBC
 1188 degrade more substantially. In contrast, DD and DP yield considerably lower win rates across all
 1189 configurations, underscoring their limited adaptability in multi-opponent settings.

Env	Opp.	RL methods		IL methods
		IDQL	OMAR	DDGIL
Push	E	0.81 ± 0.04	0.80 ± 0.04	0.81 ± 0.01
	M	0.86 ± 0.03	0.84 ± 0.05	0.82 ± 0.02
	W	0.83 ± 0.06	0.77 ± 0.04	0.84 ± 0.05
Tag	E	0.33 ± 0.08	0.33 ± 0.05	0.34 ± 0.08
	M	0.45 ± 0.03	0.39 ± 0.06	0.45 ± 0.03
	W	0.60 ± 0.07	0.58 ± 0.04	0.57 ± 0.06
Spread	E	-10.68 ± 0.25	-12.54 ± 0.06	-11.62 ± 0.41
	M	-11.91 ± 0.42	-13.93 ± 0.08	-11.87 ± 0.17
	W	-15.67 ± 0.33	-15.82 ± 0.10	-17.79 ± 0.47
Connect4	E	0.31 ± 0.02	0.32 ± 0.04	0.26 ± 0.04
	M	0.50 ± 0.03	0.47 ± 0.04	0.41 ± 0.06
	W	0.54 ± 0.02	0.53 ± 0.03	0.47 ± 0.06
Hold'em	E	0.57 ± 0.04	0.53 ± 0.02	0.55 ± 0.02
	M	0.61 ± 0.03	0.67 ± 0.01	0.62 ± 0.02
	W	0.83 ± 0.07	0.85 ± 0.03	0.88 ± 0.04
Tennis	E	0.85 ± 0.02	0.84 ± 0.04	0.81 ± 0.05
	M	0.89 ± 0.05	0.94 ± 0.06	0.90 ± 0.04
Boxing	E	0.53 ± 0.05	0.47 ± 0.06	0.48 ± 0.03
	M	0.66 ± 0.04	0.65 ± 0.07	0.52 ± 0.04

Table 5: Comparison between RL methods (IDQL, OMAR) and IL method (DDGIL). Results show that DDGIL outperforms IDQL and OMAR in some environments.

Opp. Num	Baseline	Win rate	Time / step	Memory
3	DDGIL	0.34 ± 0.08	61.3 ms	1.16 GB
	BC	0.31 ± 0.05	4.0 ms	0.31 GB
	DBC	0.38 ± 0.13	4.1 ms	0.33 GB
	DD	0.15 ± 0.03	46.4 ms	0.45 GB
	DP	0.15 ± 0.03	41.1 ms	0.45 GB
6	DDGIL	0.37 ± 0.06	72.7 ms	1.16 GB
	BC	0.32 ± 0.09	4.5 ms	0.31 GB
	DBC	0.32 ± 0.07	4.5 ms	0.33 GB
	DD	0.13 ± 0.05	51.2 ms	0.45 GB
	DP	0.15 ± 0.07	48.9 ms	0.45 GB
10	DDGIL	0.35 ± 0.05	95.2 ms	1.54 GB
	BC	0.28 ± 0.04	6.4 ms	0.33 GB
	DBC	0.31 ± 0.05	6.3 ms	0.35 GB
	DD	0.13 ± 0.08	57.1 ms	0.47 GB
	DP	0.12 ± 0.08	56.8 ms	0.46 GB

Table 6: Win rate and computational cost with increasing opponent counts in MPE-Tag.

In terms of computational cost, DDGIL incurs a higher inference time and memory usage due to the additional opponent modeling. For example, the inference time increases from 121.3 ms with 3 opponents to 251.2 ms with 10 opponents, and the memory usage increases from 1.16 GB to 1.54 GB. Although these values are larger than BC and DBC, they remain within the range of other diffusion-based baselines (DD, DP).

Overall, DDGIL demonstrates better performance scalability than BC, DBC, DD, and DP as the number of opponents increases, while its computational cost increases predictably with task complexity. This indicates that the dynamic guidance mechanism effectively stabilizes learning under more challenging multi-opponent environments.

D.6 NECESSITY OF MULTIPLE DIFFUSION MODELS IN DDGIL

In the original design, each agent is assigned an independent diffusion model, with separate conditional generation and dynamic guidance. To examine the necessity of this multi-model setup, we conducted an ablation by replacing it with a single shared model. The shared model takes the primary agent’s observation as input and outputs actions for all agents through a shared encoder and multi-head denoisers, while retaining dynamic guidance during inference. Evaluation metrics include average return, stability, and computational cost.

Env	Reward/WR	Origin Reward/WR	Time	Origin Time	Memory	Origin Mem.
Push	0.13 ± 0.03	0.81 ± 0.02	39.3 ms	47.1 ms	0.48 ± 0.00	0.59 ± 0.08
Tag	0.05 ± 0.10	0.33 ± 0.08	48.4 ms	61.3 ms	0.56 ± 0.13	0.94 ± 0.11
Spread	-22.68 ± 0.88	-11.52 ± 0.41	41.6 ms	62.4 ms	0.52 ± 0.07	0.65 ± 0.07
Reference	-38.82 ± 1.07	-26.62 ± 1.07	35.2 ms	46.2 ms	0.50 ± 0.08	0.58 ± 0.12
Tennis	0.46 ± 0.12	0.81 ± 0.05	36.5 ms	42.2 ms	0.59 ± 0.07	0.59 ± 0.01
Boxing	0.11 ± 0.09	0.47 ± 0.03	35.8 ms	44.5 ms	0.51 ± 0.09	0.60 ± 0.05
Connect4	0.15 ± 0.11	0.26 ± 0.04	37.1 ms	45.2 ms	0.47 ± 0.05	0.57 ± 0.08
Hold'em	0.34 ± 0.42	0.55 ± 0.02	33.0 ms	45.3 ms	0.47 ± 0.02	0.57 ± 0.04

Table 7: When DDGIL is implemented as a single shared model, compared with the original multi-model setup. *Origin* denotes the performance of the original model configuration. **WR** is the abbreviation for *Win Rate*.

Table 7 shows that the shared model reduces computational cost, achieving on average 26% shorter inference time (e.g., 48.4 ms vs. 61.3 ms in Tag) and about 17% lower memory usage. However, this efficiency gain comes with performance degradation in most environments. In Push, Tag, Tennis, and Boxing, win rates decline notably, while in Spread and Reference, rewards drop by more than 10 points. In contrast, Connect4 and Hold'em show smaller differences, indicating less sensitivity to model sharing.

The performance gap arises because, in the multi-model setting, the opponent model serves only as an external reference for condition and does not affect the gradient updates of the primary agent. In the shared architecture, all agents are generated by the same model, causing gradient interference, reducing the distinctiveness of guidance vectors, and hindering convergence. In addition, changes in the denoising path primary agent indirectly alter the output of other agents, breaking the design principle of using separate models to capture interaction-specific semantics through dynamic guidance.

Therefore, although the multi-model design increases computational cost, it plays a critical role in maintaining stability and controllability in imitation learning. It also preserves separable representations of the primary agent and opponent policies, thereby enhancing the interpretability and responsiveness of dynamic guidance.

E MODEL ARCHITECTURE

E.1 EXPERT REINFORCEMENT LEARNING MODEL

We configure an expert strategy model for each environment to ensure the quality and consistency of demonstration data used during the imitation learning phase. These expert models are designed to fully capture the rules and decision-making structure of the corresponding environment and remain fixed during data collection, without being updated jointly with the student policy. The architecture and training protocol of each expert differ according to the nature of the environment, as described below:

- **Atari (Tennis, Boxing):** We adopt the *Multi-Agent PPO (MAPPO)* (Logan Engstrom, 2020; John Schulman, 2017) implementation from *CleanRL* (Shengyi Huang, 2022), integrated with preprocessing utilities provided by Supersuit (Justin K. Terry, 2020a). Observations from PettingZoo are cropped, normalized, and stacked into multi-channel frames. The policy network simultaneously processes both agents' observations and outputs their respective action distributions, enabling stable learning of pixel-level competitive behavior.
- **MPE (Tag, Push, Spread, Reference):** We employ the *MADDPG* (Ryan Lowe, 2017) algorithm provided by *AgileRL* (Ustaran-Anderegg et al.), which uses a centralized Q-critic to evaluate joint action values and decentralized actors for each agent. This design effectively captures both cooperative and adversarial patterns in the multi-agent particle environment.
- **Connect4:** The expert is instantiated using the *AgileRL DQN* (Volodymyr Mnih, 2013) model with publicly available pretrained weights. These weights are obtained through

1296 curriculum learning and self-play, enabling the generation of competent gameplay demon-
 1297 strations without additional training cost.

1298

- 1299 • **Texas Hold’em:** For this imperfect-information game, we use the *Neural Fictitious Self-*
 1300 *Play (NFSP)* (Johannes Heinrich, 2016) implementation from *RLCard* (Zha et al., 2020).
 1301 NFSP maintains both a best-response policy and an average strategy memory, progressively
 1302 converging to a Nash equilibrium through self-play. This allows the expert to model strategic
 1303 inference over hidden information.
- 1304 • **Badminton:** The expert model is based on *RallyNet* (Kuang-Da Wang, 2024a), a pretrained
 1305 imitation learning model derived from real-world badminton match footage. It is capable
 1306 of predicting high-quality shot sequences and footwork trajectories, providing fluent and
 1307 realistic expert demonstrations.

1308 During the expert training phase, we log the model weights across training epochs along with the
 1309 corresponding evaluation rewards. These metrics serve as the basis for ranking and selecting expert
 1310 strengths (e.g., medium, weak) for future ablation studies. The weight checkpoint that achieves
 1311 the highest evaluation reward, typically the one saved in the final epoch, is selected as the expert
 1312 policy. All subsequent rollout datasets used for offline imitation learning are generated via interaction
 1313 between this selected expert and the environment.

1314

1315 E.2 BASELINE MODEL

1316 To benchmark the proposed framework under the offline imitation learning setting, we construct a
 1317 set of baseline models. Our survey of existing literature indicates that most multi-agent imitation
 1318 learning methods are designed for online training and are thus incompatible with our offline setting.
 1319 We therefore adopt single-agent offline imitation learning algorithms as the primary baselines, and
 1320 additionally include two multi-agent offline reinforcement learning algorithms as supplementary
 1321 experiments in Appendix D.5.

1322

- 1323 • **BC:** A classical behavior cloning model that employs a three-layer MLP with ReLU activa-
 1324 tions. It directly learns to map states to actions in a single-stage training procedure.
- 1325 • **DBC:** Based on the Diffusion Model-Augmented Behavioral Cloning (Shang-Fu Chen,
 1326 2024), this model includes a two-stage training process. A diffusion model is trained to learn
 1327 improved representations, followed by a behavior cloning policy. The decision component
 1328 is identical to the BC architecture.
- 1329 • **DP:** Implemented using the *Clean Diffuser* (Zibin Dong, 2024), Diffusion Policy
 1330 (Tim Pearce, 2023) adopts a denoising diffusion probabilistic model (DDPM). A
 1331 DIT1D-based UNet is used as the diffusion backbone, while MLPCCondition is applied
 1332 for conditional state input. The model predicts a one-step action conditioned on a short state
 1333 sequence. Originally designed for single-agent decision making, we extend it to the multi-
 1334 agent setting by conditioning on the primary agent’s state and generating action distributions
 1335 for all agents, from which the primary agent’s action is taken, like a shared-model design
 1336 with multi-head outputs.
- 1337 • **DD:** We adapt Decision Diffuser (Anurag Ajay, 2023) into an offline imitation learning
 1338 formulation using the *Clean Diffuser*. In contrast to its original reinforcement learning
 1339 design, our implementation removes reward-based conditioning and relies solely on state
 1340 information. Since each state transition in a multi-agent environment is influenced by all
 1341 agents’ actions, modeling individual-agent trajectories in isolation is insufficient. To address
 1342 this, we introduce an inverse dynamics module that predicts the actions of all agents given
 1343 consecutive states. The inverse model outputs a vector of action dimension $\text{action} \times \text{number}$
 1344 of agents, enabling accurate recovery of interaction patterns across agents.
- 1345 • **IDQL:** IDQL (Philippe Hansen-Estruch, 2023) adopts a generalized IQL architecture con-
 1346 sisting of a Q-function (critic) trained solely on dataset actions and a diffusion-parameterized
 1347 behavior policy (actor). The actor generates samples from the diffusion model and applies
 1348 importance sampling with weights computed from the critic to obtain the final policy.
 1349 Our implementation is based on the *Clean Diffuser* to support training and inference of
 1350 diffusion-based behavior policies.

	Method	Param	Push	Tag	Spread	Reference	Connect4	Hold'em	Tennis	Boxing	Badminton
1350	BC	in dim	[8, 19]	[14, 16]	18	21	64	72	64	64	17
1351		out dim	5	5	5	50	7	4	18	18	15
1352		hid dim	256	256	256	256	256	256	256	256	256
1353											
1354	DBC	Diffusion									
1355		in dim	[13, 24]	[19, 21]	23	[71, 71]	71	76	82	82	32
1356		out dim	[13, 24]	[19, 21]	23	[71, 71]	71	76	82	82	32
1357		hid dim	256	256	256	256	256	256	256	256	256
1358											
1359		BC									
1360		in dim	[8, 19]	[14, 16]	18	21	64	72	64	64	17
1361		out dim	5	5	5	50	7	4	18	18	15
1362		hid dim	256	256	256	256	256	256	256	256	256
1363											
1364	DP	in dim	[8, 19]×2	[14, 16]×2	18×2	21×2	64×2	72×2	64×2	64×2	17×2
1365		out dim	5	5	5	50	7	4	18	18	15
1366		hid dim	384	384	384	384	384	384	384	384	384
1367		step	15	15	15	15	15	15	20	20	20
1368		ex step	5	5	5	5	5	5	8	8	5
1369											
1370	DD	Diffusion									
1371		in dim	[8, 19]×H	[14, 16]×H	18×H	21×H	64×H	72×H	64×H	64×H	17×H
1372		out dim	[8, 19]×H	[14, 16]×H	18×H	21×H	64×H	72×H	64×H	64×H	17×H
1373		hid dim	320	320	320	320	320	320	320	320	320
1374		step	20	20	20	20	20	20	25	25	25
1375											
1376	InvDyn	in dim	[8, 19]×2	[14, 16]×2	18×2	21×2	64×2	72×2	64×2	64×2	17×2
1377		out dim	5×n	5×n	5×n	50×n	7×n	4×n	18×n	18×n	15×n
1378		hid dim	512	512	512	512	512	512	512	512	512
1379	IDQL	in dim	[8, 19]	[14, 16]	18	21	64	72	64	64	17
1380		out dim	5×n	5×n	5×n	50×n	7×n	4×n	18×n	18×n	15×n
1381		hid dim	512	512	512	512	512	512	512	512	512
1382	OMAR	in dim	[8, 19]	[14, 16]	18	21	64	72	64	64	17
1383		out dim	5	5	5	50	7	4	18	18	15
1384		hid dim	256	256	256	256	256	256	256	256	256
1385											
1386	DDGIL	EM									
1387		in dim	-	-	-	-	7×6×2	-	84×84×6	84×84×6	-
1388		emb dim	-	-	-	-	64	-	64	64	-
1389		hid dim	-	-	-	-	256	-	256	256	-
1390											
1391		Diffusion									
1392		in dim	[8, 19]	[14, 16]	18	21	72	64	64	64	17
1393		out dim	5	5	5	50	7	4	18	18	15
1394		hid dim	256	256	256	256	256	256	256	256	256
1395		step	15	15	15	15	15	15	20	20	20
1396											
1397											

Table 8: Model Architecture Parameters for baseline. Hyperparameter configurations for all baseline models and our proposed DDGI across seven environments.

- **OMAR**: Offline Multi-Agent RL with Actor Rectification (OMAR) (Ling Pan, 2022) combines first-order policy gradients with zeroth-order optimization to address the non-concavity of conservative value functions in the actor parameter space, reducing the risk of suboptimal convergence. This design mitigates global coordination failures caused by suboptimal policies of individual agents in offline multi-agent reinforcement learning. Our implementation follows the original OMAR architecture to support policy optimization in multi-agent control tasks.

The hyperparameter settings for all baseline models are summarized in Table 8. For clarity, we define the following abbreviations: **in dim** (input dimension), **hid dim** (hidden dimension), **out dim** (output dimension), **step** (diffusion sampling steps), **ex step** (extra steps in DP), **EM** (embedding model in DDGIL), and **InvDyn** (inverse dynamics module in DD). In DD, **H** is the trajectory horizon (defaulting to the environment’s episode length), while **n** denotes the number of agents. For Push/Tag, state dimensions are shown as (agent, opponent). In Tennis, Boxing, and Connect4, state dimensions follow DDGIL’s embedding size.

1404	Method	Param	Push	Tag	Spread	Reference	Connect4	Hold'em	Tennis	Boxing	Badminton
1405	-	bs	256	256	256	256	256	256	64	64	256
1406	BC	lr	1e-4	1e-4	1e-4	1e-4	1e-4	1e-4	1e-4	1e-4	1e-4
1407		epoch	300	300	300	300	300	300	400	400	300
1408	DBC	Diffusion									
1409		lr	5e-4	5e-4	5e-4	5e-4	5e-4	5e-4	1e-5	1e-5	5e-4
1410		epoch	1000	1000	1000	1000	1000	1000	1000	1000	1000
1411		BC									
1412		lr	5e-4	5e-4	5e-4	5e-4	5e-4	5e-4	1e-5	1e-5	5e-4
1413		epoch	300	300	300	300	300	300	400	400	300
1414	DP										
1415		lr	5e-4	5e-4	5e-4	5e-4	5e-4	5e-4	1e-4	1e-4	5e-4
1416		epoch	1000	1000	1000	1000	1000	1000	1200	1200	1000
1417	DD	Diffusion									
1418		lr	5e-4	5e-4	5e-4	5e-4	5e-4	5e-4	1e-5	1e-5	5e-4
1419		epoch	5000	5000	5000	5000	5000	5000	8000	8000	6000
1420		InvDyn									
1421		lr	5e-4	5e-4	5e-4	5e-4	5e-4	5e-4	1e-5	1e-5	5e-4
1422		epoch	2000	2000	2000	2000	2000	2000	3000	3000	2000
1423	DQL										
1424		lr	5e-4	5e-4	5e-4	5e-4	5e-4	5e-4	5e-4	5e-4	5e-4
1425		epoch	1000	1000	1000	1000	1000	1000	1000	1000	1000
1426	OMAR										
1427		lr	1e-5	1e-5	1e-5	1e-5	1e-5	1e-5	1e-5	1e-5	1e-5
1428		epoch	5000	5000	5000	5000	5000	5000	5000	5000	5000
1429	DDGIL	EM									
1430		lr	-	-	-	-	5e-4	-	1e-5	1e-5	-
1431		epoch	-	-	-	-	100	-	100	100	-
1432		Diffusion									
1433		lr	1e-5	1e-5	1e-5	1e-5	1e-5	1e-5	1e-5	1e-5	1e-5
1434		epoch	500	500	500	500	500	500	600	600	500

Table 9: Model Training Parameters for baseline in each environment.

Table 9 summarizes the hyperparameter configurations for all baseline models. The following abbreviations are used: **lr** (learning rate), **bs** (batch size), and **epoch** (total training epochs). All baselines are optimized using Adam. The batch size for each environment is listed at the top of the table.

E.3 DDGIL MODEL

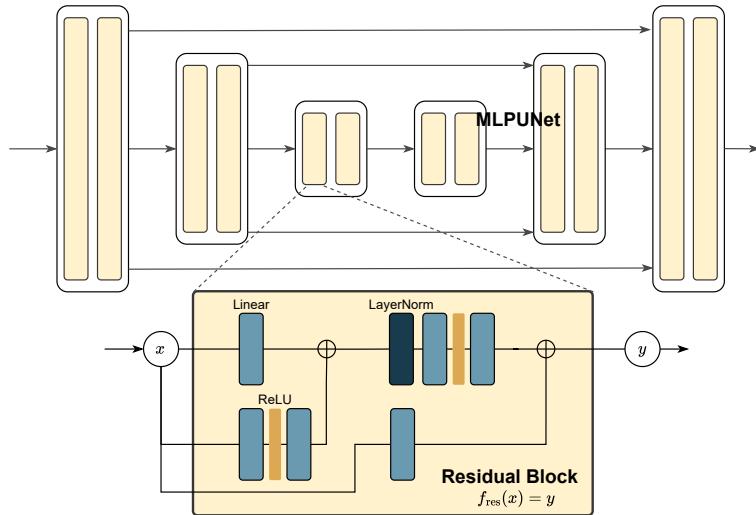
To handle the diversity of state representations across environments, we structure DDGIL into two modular yet interdependent components: an embedding module and a diffusion-based policy generation module. The former transforms high-dimensional, non-vector observations into compact latent embeddings, while the latter serves as the core mechanism for skill-conditioned policy generation.

Embedding Module. The embedding module is responsible for preprocessing non-vectorial state inputs, such as images or structured tensors, and converting them into fixed-length latent representations suitable for conditioning the diffusion model. The training algorithm can be found in Appendix B. We implement two variants of the embedding module tailored to different input types:

- For **Atari** (e.g., Boxing, Tennis), the encoder consists of three Conv2D layers, and the decoder mirrors this with three ConvTranspose2D layers. The raw input of shape $84 \times 84 \times 6$ is normalized before being encoded into a 64-dimensional latent vector.
- For **Connect4** The state is a $7 \times 6 \times 2$ one-hot tensor, we adopt a lightweight MLP encoder comprising two linear layers with nonlinear activations, along with a symmetric two-layer MLP decoder. No normalization is applied to the input, as it is already structured and non-visual.

1458 For vector-based environments such as MPE and Texas Hold’em, no embedding is required; the
 1459 original state vectors are passed directly to the diffusion model.
 1460

1461 **Diffusion Architecture.** The policy generation component is implemented using a modified
 1462 MLPUnet architecture, inspired by the UNet1D structure (Michael Janner, 2022). Our version
 1463 retains LayerNorm within the MLP layers for training stability, removes the original attention
 1464 modules to reduce computational overhead, and replaces the conditional input mechanism (ConID)
 1465 with an MLP-based residual block. This design simplifies the conditioning pathway while preserving
 1466 expressivity. The architecture is illustrated in Figure 7.
 1467
 1468



1469
 1470
 1471
 1472
 1473
 1474
 1475
 1476
 1477
 1478
 1479
 1480
 1481
 1482
 1483
 1484
 1485
 1486
 1487 Figure 7: UNet Model architecture in DDGIL. We adopt an MLP module in our DDGIL model, without
 1488 incorporating any attention modules.
 1489
 1490

E.4 COMPUTING RESOURCES

1491 All models are trained and evaluated on an RTX 3090 GPU with 24GB of memory. Both agent
 1492 and opponent policies are pre-trained using offline datasets. During inference, the diffusion model
 1493 performs a denoising sampling loop, with the number of steps adjusted based on task complexity.
 1494

1495 **Training Cost.** To illustrate computational cost, we report training times for two representative
 1496 environments: MPE-Push and Atari-Boxing. Training an MPE model takes approximately 10 minutes,
 1497 while Atari and Badminton require longer due to larger state-action spaces. A summary of these
 1498 results is provided in Table 10. All experiments are conducted under consistent hardware settings to
 1499 ensure reproducibility.
 1500

Environment	Dataset Size	Time Cost
MPE-Push	50	3m 16s
	100	10m 34s
	250	18m 04s
	500	25m 26s
Atari-Boxing	10	2m 36s
	50	8m 42s
	100	18m 31s
	250	35m 14s

1501
 1502
 1503
 1504
 1505
 1506
 1507
 1508
 1509
 1510
 1511 Table 10: Training time by dataset size.

Inference Cost. We evaluate per-step inference time and memory usage under a fixed setting of 100 episodes against an expert opponent, averaged over three seeds with standard deviation, as shown in Table 11. Compared to BC and DBC, which require only a single forward pass, DDGIL performs a multi-step denoising process, resulting in longer inference time. Compared to DP and DD, DDGIL additionally applies dynamic guidance at each denoising step, further increasing computational cost. For example, in MPE-Push, it requires 67.1 ms per step, slightly higher than DP (32.3 ms) and DD (33.7 ms), primarily due to the added guidance operations within the diffusion sampling process. Memory usage is approximately 0.59 GB for DDGIL, compared to 0.45 GB for DP and DD, and is notably higher than non-diffusion baselines. Despite the increased cost, results in Section 5.3 and Section 5.4 show that DDGIL’s enhanced stability and robustness justify the trade-off, particularly under opponent diversity or strategy distribution shifts.

Environment	Time Cost	Memory Cost
MPE-Push	47.1 ms	0.59 GB
Atari-Boxing	44.5 ms	0.60 GB

Table 11: Inference cost: time per step and memory.

F INFERENCE DETAILS

To evaluate policy robustness under varying opponent strategies, we categorize adversaries into three skill levels: **Expert**, **Medium**, and **Weak**. We train a reference RL agent for most environments and record reward curves and loss values at each checkpoint. If reward logs are available, we select weights corresponding to the highest, median, and lowest rewards to represent O_{expert} , O_{medium} , and O_{weak} opponents, respectively. When reward information is unavailable, model checkpoints are chosen based on loss convergence, with the lowest-loss weights serving as the Expert baseline.

In Connect4, we use the pretrained DQN model provided by AgileRL without additional training. In addition to the DQN, AgileRL offers a rule-based agent with three modes: Strong, Weak, and Random. Empirical evaluations reveal that the Strong policy outperforms the pretrained DQN, while the DQN marginally surpasses the Weak policy. Accordingly, we designate:

- **Strong** as Expert
- **DQN** as Medium
- **Weak** as Weak

Although the Weak agent performs similarly to the Medium baseline in some settings, we retain the above classification for consistency.

In the Badminton environment, skill levels are not defined by training performance but rather by real-world players’ historical outcomes. The players and data used in this setting derive from professional matches. We select three professional players from the dataset, referred to as K, C, and V, to represent distinct playing styles in our experiments.

Additionally, we employ the RallyNet model as an opponent. RallyNet is a hierarchical offline imitation learning model that reproduces realistic stroke patterns and strategic play styles. We evaluate our method through interactions between RallyNet and other baselines across all six pairwise matchups among the three players.

G DATASET INFORMATION

G.1 DATASET CONSTRUCTION

After training baseline models for each environment in Appendix E, we generate offline datasets by interacting with the expert policies described in Appendix F. Specifically, we collect 500 trajectories for MPE environments (Push, Tag, Spread, Reference) and Classic environments (Connect4, Texas Hold’em), and 250 trajectories for Atari environments (Boxing, Tennis). For Badminton, no additional rollout is required, as the environment includes pre-collected match data for real-world players.

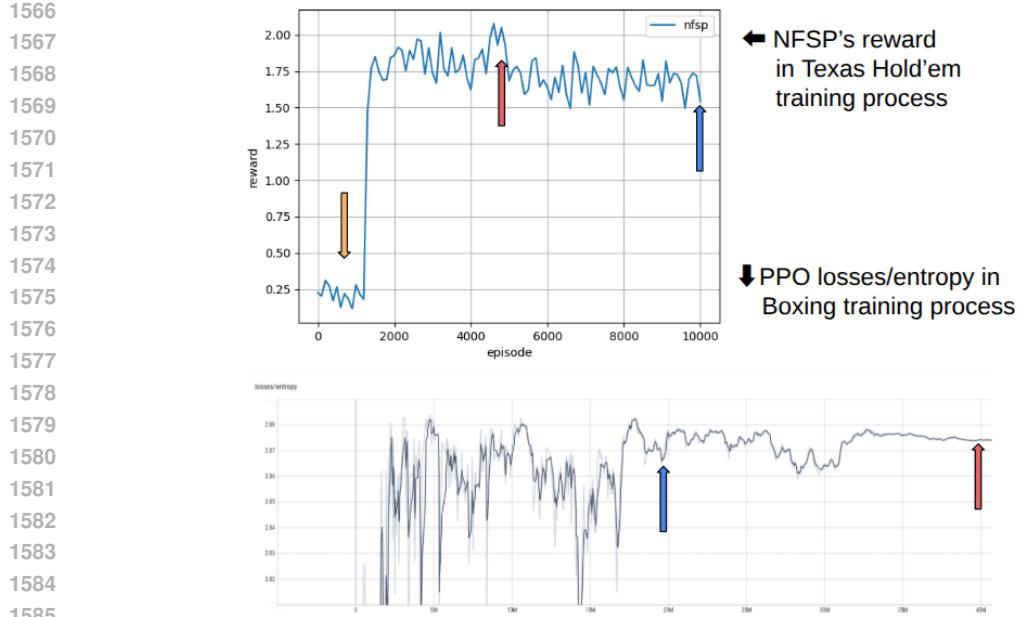


Figure 8: Example training curves used to define skill levels. The top-left plot shows the evaluation reward of NFSP in the Texas Hold'em environment. We selected the highest reward checkpoint as the Expert policy and used checkpoints around 10,000 and 1,000 rewards as Medium and Weak, respectively. The bottom plot shows the loss and entropy of PPO in the Boxing environment; we select the final stage (around 40M steps) as Expert and the mid-stage (around 20M steps) as Medium.

In offline imitation learning, we assume access to a dataset $\mathcal{D} = \{\tau_1, \tau_2, \dots, \tau_{|\mathcal{D}|}\}$. Each trajectory $\tau = \{(S_0, A_0), (S_1, A_1), \dots, (S_H, A_H)\}$ has a fixed horizon H , with $S_t = \{s_t^1, \dots, s_t^K\}$ and $A_t = \{a_t^1, \dots, a_t^K\}$ denoting the joint observations and actions of all K agents at timestep t .

For implementation, each transition can be further represented as a tuple $\{s_{t-1}^i, a_{t-1}^i, s_t^i, a_t^i\}$, for agent i and timestep $t \leq H$. At $t = 0$, we initialize $s_{-1}^{(i)}$ and $a_{-1}^{(i)}$ as zero vectors to maintain format consistency and enable models to process historical information. The full trajectory can also be viewed as overlapping segments of such tuples:

$$\{s_{t-1}^{(i)}, a_{t-1}^{(i)}, s_t^{(i)}, a_t^{(i)}\}, \quad \{s_t^{(i)}, a_t^{(i)}, s_{t+1}^{(i)}, a_{t+1}^{(i)}\}, \quad \dots$$

This flexible data structure supports different forms of training batches depending on model requirements:

- **State-action pairs:** $\{s_t^{(i)}, a_t^{(i)}\}$, for standard behavior cloning.
- **Transition tuples:** $\{s_{t-1}^{(i)}, a_{t-1}^{(i)}, s_t^{(i)}, a_t^{(i)}\}$, for models with temporal dependency (e.g., DBC).
- **Trajectory segments:** $\tau^{(i)} = \{s_0^{(i)}, a_0^{(i)}, \dots, s_n^{(i)}, a_n^{(i)}\}$, $n \in [0, H]$, for sequence-based models or diffusion samplers.

G.2 DATASET COLLECTION IN DIFFERENT TYPES OF ENVIRONMENTS

Following PettingZoo’s environment taxonomy, we organize the datasets under two interaction types: **AEC (Agent Environment Cycle)** and **Parallel**.

- **AEC environments:** At each time step t , only one agent receives an observation and acts. The state s_t stores the active agent’s observation, while the action a_t is constructed by concatenating the agent’s action at t with the opponent’s action at $t-1$.

1620
 1621 • **Parallel environments:** All agents observe and act simultaneously at each time step. The
 1622 full state s_t is a concatenation of all agent observations $\{s_t^i\}$, and a_t is a concatenation of
 1623 all actions $\{a_t^i\}$ in a fixed environment-defined order. Maintaining this order is crucial to
 1624 ensure proper alignment during training.

1625 During baseline training, we split the dataset by agent name and ordering so that each model
 1626 only learns from its data. This format is designed primarily to support DBC, which requires
 1627 $\{s_{t-1}, a_{t-1}, s_t\}$ as input to predict a_t . Other baselines use only $\{s_t, a_t\}$ pairs, making this unified
 1628 structure compatible across methods.

1629 Although no environment rollout is performed for Badminton, the raw match data must be filtered
 1630 and reformatted. We select a fixed subset of columns as state and action features, structuring the data
 1631 as $\{s_{t-1}, a_{t-1}, s_t, a_t\}$. Both states and actions contain discrete and continuous variables; discrete
 1632 features are one-hot encoded and concatenated with continuous values.

1633 Specifically, the state vector includes 1 discrete feature (classes 0–10) and 6 continuous features,
 1634 resulting in 17 dimensions after encoding. The action vector comprises 1 discrete feature (classes
 1635 0–10) and 4 continuous elements, totaling 15 dimensions. These hybrid vectors are used for training
 1636 after preprocessing.

1638 H THE USE OF LARGE LANGUAGE MODELS (LLMs)

1640 During manuscript preparation, we used OpenAI’s ChatGPT (GPT-5) to refine wording and improve
 1641 clarity. The model was occasionally consulted for alternative perspectives in method design and for
 1642 assistance in presenting mathematical derivations more coherently. We also employed Grammarly for
 1643 grammar and style checking.

1644 All methodological contributions, theoretical developments, algorithmic designs, and empirical
 1645 analyses were conceived and validated by the authors, and these tools served only as auxiliary aids
 1646 for writing and presentation.

1648 I LIMITATIONS AND DISCUSSION

1651 While DDGIL achieves strong performance across diverse benchmarks, several limitations are noted.
 1652 The inference procedure requires recomputing dynamic weights at each denoising step, improving
 1653 adaptability but incurring higher latency and memory usage than fixed-weight policies. Future work
 1654 may consider reducing denoising steps or distilling the sampler into a lighter model.

1655 Another limitation is the need to train a separate diffusion model for each agent. This ensures
 1656 stability by modeling heterogeneous behaviors but scales poorly as agent numbers grow (see ablation
 1657 experiment D.5). Future work may consider parameter-efficient designs, such as shared encoders
 1658 with agent-specific decoders, to improve scalability.

1659 The discrepancy signal for dynamic weighting is computed by averaging opponent predictions. While
 1660 suitable in our benchmarks with relatively homogeneous opponents, this may be less effective under
 1661 heterogeneous or adversarial settings. Future work may consider weighted aggregation or attention
 1662 mechanisms for more expressive signals.

1663 Although we did not explicitly consider $w_D < 0.5$, fixed-weight ablations (see experiment D.2) show
 1664 that performance degrades when weighting is closer to the unconditional update, suggesting that
 1665 bounding w_D away from the unconditional case aids stability.

1666 Our evaluation does not cover larger multi-agent benchmarks such as MaMuJoCo and SMAC,
 1667 although the formulation of DDGIL does not rely on benchmark-specific assumptions and may
 1668 extend to such domains. Further empirical validation remains an important direction for future work.

1670 Finally, adaptation in DDGIL is confined to inference and was not tested under severe distributional
 1671 shifts. Future work may extend the framework with lightweight online adaptation or meta-learning to
 1672 improve generalization in such regimes.