Hierarchical Policy-Gradient Reinforcement Learning for Multi-Agent Shepherding Control of Non-Cohesive Targets

Stefano Covone[1], Italo Napolitano[1], Francesco De Lellis[2], Mario di Bernardo[1,2,*]

Abstract— We propose a decentralized reinforcement learning solution for multi-agent shepherding of non-cohesive targets using policy-gradient methods. Our architecture integrates target-selection with target-driving through Proximal Policy Optimization, overcoming discrete-action constraints of previous Deep Q-Network approaches and enabling smoother agent trajectories. This model-free framework effectively solves the shepherding problem without prior dynamics knowledge. Experiments demonstrate our method's effectiveness and scalability with increased target numbers and limited sensing capabilities.

I. INTRODUCTION

The shepherding problem in robotics exemplifies the problem of harnessing complex systems for control [1], [2]. It generally involves a group of actively controlled agents, termed *herders*, strategically influencing a group of passive agents, referred to as *targets*. From a control-theoretic perspective, this represents a paradigmatic example of indirect control in complex systems [3], where an active system manipulates a passive one to guide its collective behavior toward a desired state.

This problem has wide-ranging applications: from wildlife management using UAVs to prevent animals from entering dangerous areas [4], to security tasks in search and rescue operations [5] and protection of strategic areas [6].

State-of-the-art shepherding solutions face several challenges. Model-based approaches [7] create computationally intensive optimization problems that scale poorly, while heuristic methods [8], [9] typically use predefined two-phase strategies where herders first collect targets and then drive them toward a desired region.

In contrast, learning-based approaches have emerged as powerful alternatives for shepherding, enabling agents to develop effective strategies without predefined rules or prior knowledge of target dynamics. For example, in [10], herders learn to steer flocks in complex structured environments via Deep Q-Networks. Deep policy-gradient Reinforcement Learning (RL) has further advanced these capabilities. In [11], separate driving and collecting policies are learnt

using threshold-based switching, while curriculum learning is employed in [12] for driving and mode-selection. These methods, however, are limited to single-herder scenarios.

For multi-herder coordination, Multi-Agent Reinforcement Learning (MARL) offers promising solutions [13]. Deep Deterministic Policy Gradient (DDPG) is applied in [14] for centralized learning of driving and collecting behaviors, while Proximal Policy Optimization (PPO) with convolutional networks is used in [15] to develop scalable, decentralized policies.

Despite their impressive performance, these approaches rely on the simplifying assumption that targets behave as a cohesive flock [8] – a hypothesis that proves unrealistic in various domains including emergency evacuation scenarios, wildlife management of non-herding species, and heterogeneous robotic swarms [1]. Without cohesion, as noted in [7], the problem complexity increases significantly as herders must control each target individually rather than influencing the group's center of mass.

To address this scenario, herder behavior is divided in [16] into two distinct tasks: *target-selection*, where herders decide what targets to chase, and *driving*, where herders push the selected targets towards the goal region. Another approach [17] uses Q-learning with heuristic rules, while [18] employs DQN but suffers from discrete action limitations.

In this work, we leverage Proximal Policy Optimization (PPO), a state-of-the-art deep policy-gradient algorithm [19], aiming to: (i) operate within a fully learning-based framework, overcoming the reliance on heuristic components found in hybrid approaches, (ii) require no prior knowledge of target dynamics, and (iii) achieve state-of-the-art performance by utilizing a continuous action space, rather than a discrete one as in DQN.

II. PROBLEM STATEMENT

Consider a square domain $\mathcal{D} = [-R, R]^2 \subset \mathbb{R}^2$, where N herders must guide M targets to a circular goal region $\Omega_G \in \mathcal{D}$, centered at the origin, with radius $\rho_G < R$. Initially, both the herders and the targets are uniformly randomly distributed within a circle $\Omega_0 \in \mathcal{D}$ of radius R, centered at the origin.

Let $\mathbf{H} \in \mathbb{R}^{2N}$ denote the vector of herder positions, defined as $\mathbf{H} = [\mathbf{H}_1, \ldots, \mathbf{H}_N]$, where $\mathbf{H}_i \in \mathbb{R}^2$ represents the Cartesian coordinates of the i-th herder. Similarly, let $\mathbf{T} \in \mathbb{R}^{2M}$ denote the vector of target positions, defined as $\mathbf{T} = [\mathbf{T}_1, \ldots, \mathbf{T}_M]$, where $\mathbf{T}_a \in \mathbb{R}^2$ represents the Cartesian coordinates of the a-th target.

Fig. 1: Two-layer feedback control scheme: each herder $\mathbf{H}_{i,j}$ detects the other agents' positions and selects the target $\mathbf{T}_{i,j}^*$ to control via the *target-selection* policy, which is then driven according to the *driving* policy, that outputs the velocity $\mathbf{u}$ of the corresponding herder.

Building on the minimal shepherding model presented in [1], both populations are modeled with first-order differential equations under the assumption of a negligible acceleration phase duration [20], providing sufficient dynamic fidelity while maintaining computational efficiency for the learning framework. Unlike most existing studies, this model does not assume cohesive collective behavior among the targets; instead, their dynamics are governed by the overdamped Langevin equation:

$$\dot{\mathbf{T}}_a = \sqrt{2D}\boldsymbol{\eta} + k^{\mathrm{T}} \sum_{i=1}^{N} \boldsymbol{\Theta}(\lambda, \mathbf{T}_a - \mathbf{H}_i) \qquad (1)$$

where $\boldsymbol{\eta} \in \mathbb{R}^2$ represents standard Gaussian noise and $\boldsymbol{\Theta}(\lambda, \mathbf{x})$ defines an inter-population interaction of the form

$$\boldsymbol{\Theta}(\lambda, \mathbf{x}) = \begin{cases} (\lambda - |\mathbf{x}|)\hat{\mathbf{x}} & |\mathbf{x}| \leq \lambda \\ 0 & |\mathbf{x}| > \lambda \end{cases}, \qquad (2)$$

representing a harmonic repulsion exerted by all nearby herders on the a-th target, which is typically adopted in Physics as the simplest type of repulsive force among particles [21]. Specifically, herders repel targets within a characteristic distance λ, with intensity regulated by k^{T} [1]. The maximum escape speed of a target under repulsion from a nearby herder is given by $v_{\mathrm{T}} = k^{\mathrm{T}}\lambda$, assuming $k^{\mathrm{T}}\lambda^2 \gg D$ to ensure that the repulsive force dominates the targets' Brownian motion.

The herders are modeled as single integrators following the control law – e.g. mobile robots with a feedback linearized dynamics – as commonly used in the robotics literature for this problem [22]:

$$\dot{\mathbf{H}}_i = \mathbf{u}_i \qquad (3)$$

where $\mathbf{u}_i \in \mathbb{R}^2$ represents the control input for the i-th herder, with a maximum velocity $v_{\mathrm{H}} > k^{\mathrm{T}}\lambda$ to prevent the formation of stable chasing patterns [1]. The model is simulated using the Euler–Maruyama integration method. The nominal parameters, provided in the Appendix, are taken from [1], except for the diffusion coefficient D, which is set to a smaller value.

The goal is to design a control law $\mathbf{u}_i$ for $i = 1, \ldots, N$ such that every target is driven and contained within the designated goal region Ω_{G} in finite time. A buffered region $\hat{\Omega}_G$ with radius $(1 + \varepsilon)\rho_G$ is introduced to accommodate for stochastic effects. We consider a decentralized, communication-free setting where herders rely solely on their observations, and can sense agents in the entire domain but do not exchange internal information, i.e., their decisions, as considered in [23].

III. CONTROL ARCHITECTURE

To manage the complexity of the task, we adopt a two-layer hierarchical approach, as also done in [18]. As illustrated in Fig. 1, a high-level decision-making policy manages the *target selection* subtask, while a low-level control policy governs the *target driving* subtask. This task decomposition aligns with existing methods in the literature on shepherding strategies for non-cohesive targets [16], [24]. Each layer is trained in a specific herding scenario: the low-level policy is learned in a single-herder, single-target setting, while the high-level policy is trained in a multiple-herder, multiple-target environment with the previously learned low-level policy kept fixed. The resulting hierarchical control architecture enables each herder to independently select a target using the high-level policy and then interact with it effectively through the guidance of the specialized low-level policy.

In what follows, we leverage a continuous action space with PPO [19] to learn both low-level and high-level policies, overcoming the limitations of previous approaches such as [18]. We extend traditional PPO to cooperative multi-agent settings using Multi-Agent Proximal Policy Optimization (MAPPO) [25]. To rigorously validate our approach, we compare our trained policy against the state-of-the-art model-based strategy proposed in [1].

A. Metrics

For validation we define $\chi(k)$ as the fraction of targets inside the buffered goal region at time step k:

$$\chi(k) = \frac{|\{a : \mathbf{T}_a(k) \in \hat{\Omega}_{\mathrm{G}}, \, a \in [1, M]\}|}{M}, \qquad (4)$$

where $|\,\mathcal{A}\,|$ denotes a set $\mathcal{A}$ cardinality. The settling time $n^\star$ is the first time step when all targets enter and remain in $\hat{\Omega}_{\mathrm{G}}$:

$$n^\star = \inf_n \{n \ge 0 : \chi(k) \ge 0.99, \forall k \in [n, n_{\mathrm{f}}]\}, \qquad (5)$$

where $n_{\mathrm{f}} = \min{(n + n_{\mathrm{t}}, n_{\mathrm{h}})}$. An episode terminates when all targets remain within the goal region for n_{t} consecutive steps or if the maximum step count n_{h} is reached. The problem is solved if $n^\star$ is finite.

We assess efficiency of the policy using the path length, defined as the average distance traveled by the herders over a finite time interval $[0, n]$:

$$d(n) = \frac{1}{N} \sum_{i=1}^{N} \sum_{k=0}^{n-1} \|\mathbf{H}_i(k+1) - \mathbf{H}_i(k)\|. \qquad (6)$$

In our simulations, we evaluate the path length $d = d(n_{\mathrm{f}})$, which also indicates the scaled average control effort, given the herder dynamics in (3).

IV. DRIVING SUB-TASK: ONE HERDER-ONE TARGET

The low-level policy for the driving subtask is trained in a single-herder, single-target scenario ($N = 1$, $M = 1$) with normalized target position and herder-target relative position as inputs.

While simple reward signals can theoretically produce complex behaviors, carefully designed reward functions enhance convergence and performance [26]. Our reward function (inspired from the one used in [20]) balances four key objectives for the herder: (i) it should approach the target to enter its influence zone, (ii) position itself to guide the target toward the goal, (iii) minimize control effort, and (iv) recognize goal achievement. Hence, the reward function is shaped as:

$$\begin{aligned} r_{\mathrm{D},k} = & -k_1 \|\mathbf{T}(k) - \mathbf{H}(k)\| \mathbb{1}_{\{\mathbf{T}(k) \in \Omega_{\mathrm{G}}\}} + \\ & -k_2 (\|\mathbf{T}(k)\| - \rho_G) \mathbb{1}_{\{\mathbf{T}(k) \in \Omega_{\mathrm{G}}\}} - k_3 \|\mathbf{u}(k)\| \end{aligned} \qquad (7)$$

where $\mathbb{1}_{\{x\}}$ is 1 when x is true and 0 otherwise.

As reported in the Appendix, the gains k_1, k_2, k_3 were systematically determined by their relative importance: k_2 (goal guidance) received the highest value as our primary objective, k_1 (target approach) was set to an intermediate value to bootstrap early learning, and k_3 (efficiency) was assigned the lowest value as a secondary objective. This hierarchy ($k_3 < k_1 < k_2$) creates a natural progression that guides policy development through increasingly complex behaviors.

In our implementation of PPO, both the actor and the critic neural networks feature four input neurons and five hidden layers each made of 64 neurons with ReLU activation. As output, the critic network has a single neuron with a linear activation function, while the actor network has two neurons with tanh activation functions. During training, control actions are sampled from Gaussian distributions whose means are provided by the actor's output, while their standard deviations are extra parameters of the actor network subject to the optimizer updates [27]. Training occurred


Fig. 2: Learning curves during training: **(a)** *Driving* policy cumulative reward, smoothed via moving average of 200 samples; **(b)** *Target-selection* policy cumulative reward, smoothed via moving average of 2000 samples. For both policies only the first half of the training is shown, to highlight the learning phase.

over $E = 2 \cdot 10^4$ episodes of $n_{\mathrm{h}} = 1200$ steps each, with $n_{\mathrm{t}} = 200$ steps. The training curve in Fig. 2a shows rapid initial improvement followed by stable convergence, indicating effective reward shaping.

Our PPO implementation follows [28], with hyperparameters initially based on [19] and refined through systematic experimentation to optimize sample efficiency and learning stability.

A. Validation

The policy was tested across $E = 1000$ episodes and compared to the model-based driving policy proposed in [1], where the herder steers the target by maintaining a fixed distance behind it, applying a proportional control strategy.

As illustrated in Fig. 3, our RL-driven herder agent initially approaches the target, guides it toward the goal region, and then maintains its position using minimal movements, demonstrating the effectiveness of our reward formulation.

While both strategies successfully completed the task in all the episodes, Fig. 5a shows how our RL policy significantly outperformed the heuristic strategy both in performance and efficiency.

V. MULTIPLE HERDERS - MULTIPLE TARGETS SCENARIO

In the case $N > 1, M > 1$, a MAPPO agent is trained for the target selection subtask. The learning process follows the centralized-training, decentralized-execution paradigm with parameter sharing [13]. Specifically, during training, all herders contribute to learning a shared policy by updating the same neural network parameters. Then, once trained, each herder independently exploits this policy to select a target. Due to constraints in the neural network structure on input and output dimensions, we consider the specific case of $N = 2$ herders and $M = 5$ targets. Each herder receives as input its own position, the position of the other herder,

Fig. 3: Example of the learned *driving* policy in a single herder, single target setting: the herder (blue diamond) approaches the target (magenta circle), drives it to the goal region (green circle) and contains it. Big markers show initial positions, small ones show final positions.
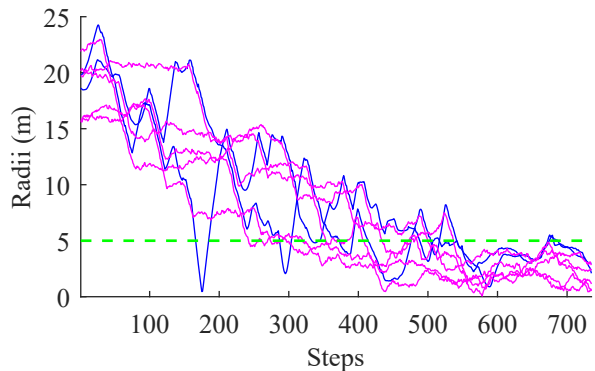

Fig. 4: Validation example of the fully learning-based solution in the $N = 2, M = 5$ scenario: the radii of the herders (blue lines) and targets (magenta lines) are shown, compared to the goal region radius $\rho_G = 5$ (green dashed line). The herders successfully steer and contain the targets in the goal region.

and the positions of all targets in the environment, and then outputs the index of the selected target. It is important to note that decision-making is decentralized, as each herder senses the positions of other agents, but selects a target independently and without communication with others.

For target selection, we define the reward function $r_{T,k}$ as:

$$r_{T,k} = -k_4 \sum_{a=1}^{M} (\|\mathbf{T}_a(k)\| - \rho_G) \mathbb{1}_{\{\mathbf{T}_a(k) \in \Omega_G\}} \qquad (8)$$

Thus, at time k, the agents receive an increasing penalty for each target located outside the goal region, scaled by $k_4 > 0$.

The MAPPO agent is implemented using an Actor-Critic neural network architecture. Both networks have $2(N + M)$ input neurons, representing the absolute positions of all the agents normalized with respect to the arena length, followed by 2 hidden layers of 256 and 128 neurons, respectively, each activated by ReLU functions. The critic outputs the state value through a single neuron, while the actor has M output neurons corresponding to the selectable targets. At each step, the outputs are normalized to sum to 1, representing the probability distribution over target selection.

The target selection policy is trained over $E = 2 \cdot 10^5$ episodes, each lasting $n_h = 3000$ steps with $n_t = 200$ steps.

To facilitate training, we leverage the modularity of our hierarchical control architecture by fixing target selection for $n_w = 100$ consecutive time steps, allowing sufficient time to influence targets and evaluate the selection's effectiveness. After training, we relax this assumption, allowing the herder to select a target at each time step.

In Fig. 2b, we show the cumulative reward curve, while the corresponding reward weights and MAPPO hyperparameters are listed in the Appendix. Hyperparameters tuning largely began from the values used in the *driving* policy, as these proved effective. Adjustments were made to the horizon and the number of actors, reflecting the reduced number of forward steps per episode. Additionally, the entropy coefficient

was set to zero, as the policy stochasticity itself provided sufficient exploration. Following [25], the use of minibatches was discarded.

A. Validation

The learned policy is validated over $E = 1000$ episodes with seeded initial conditions uniformly sampled in Ω_0. Each episode lasts $n_h = 3000$ steps and may terminate early if $n_t = 200$ successful steps are reached.

For comparison, we use the target selection law from [1]. In this approach, the i-th herder selects the target $\mathbf{T}_i^*$ furthest from the goal region, satisfying $\|\mathbf{T}_i^* - \mathbf{H}_i\| \leq \|\mathbf{T}_i^* - \mathbf{H}_j\|$, $\forall i \neq j$. This rule partitions the set of selectable targets among herders, enforcing cooperation in target selection without requiring explicit communication, but only positional information. Each target-selection strategy is combined with the corresponding driving policy.

RL agents can effectively learn decision-making strategies in a model-free fashion, as proven by Fig. 5b: both strategies successfully solve all the episodes, but the RL agent achieves better settling times and path lengths.

Fig. 4 illustrates an example episode in which two herders successfully steer five passive agents toward the goal region and contain them, as indicated by the targets' radii remaining within the goal region radius.

B. Robustness analysis

To assess the robustness of the trained policies under parametric variations, the target noise D, repulsion distance λ, and repulsion strength k^T are varied according to normal distributions, centered at the nominal values provided in the Appendix, with a standard deviation of 30%. The evaluation is conducted over $E = 1000$ validation episodes with seeded initial conditions.

Numerical results for the scenarios with $N = 1, M = 1$ and $N = 2, M = 5$ are illustrated in Fig. 5c and
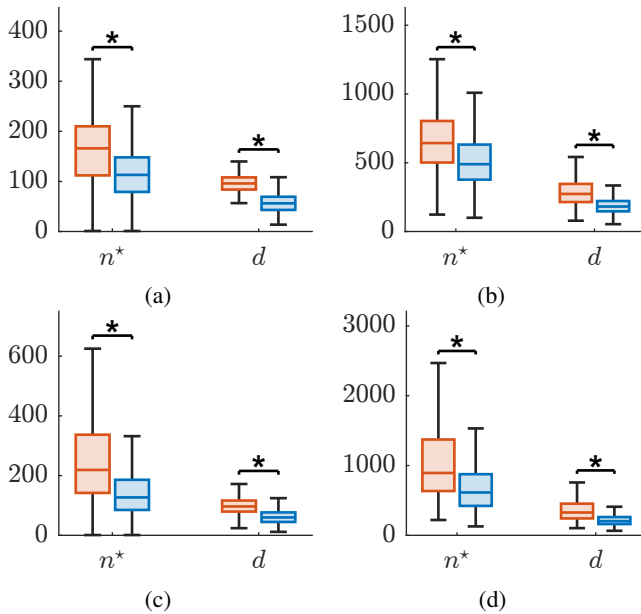
Fig. 5: Validation results for the learning-based strategies (blue) over 1000 episodes with seeded initial conditions, showing average settling time $n^\star$ in steps and average path length d in meters. Compared against the heuristic approach (orange) from [1] for both the *(a)* $N = 1$, $M = 1$ and *(b)* $N = 2$, $M = 5$ configurations. A robustness analysis is also presented by varying the targets' model parameters by 30% around their nominal values, for both *(c)* $N = 1$, $M = 1$ and *(d)* $N = 2$, $M = 5$ settings. Box plots are shown for each metric. Mann-Whitney U test was performed on each metric pair yielding p-values always smaller than 0.001.

Fig. 5d, respectively. In both scenarios, the RL-based approach demonstrates a significant improvement over the heuristic strategy, achieving superior performance in terms of both settling time and path length. Furthermore, the RL method consistently attains a success rate exceeding 97.50%, markedly higher than the heuristic approach's success rate of 84.20%. Comparable outcomes are also observed when considering a 20% variation in the system parameters, where the RL strategy again shows greater robustness relative to the heuristic method.

C. Extension to large-scale settings

Unlimited sensing capabilities for herders is often unrealistic in practical applications. To address this limitation, we implement topological sensing (c.f. [29]) where each herder perceives only the $\hat{N} \le N$ nearest herders (including itself) and the $\hat{M} \le M$ closest targets. This approach allows us to extend our target-selection policy trained for $(\hat{N}, \hat{M})$ agents to much larger environments ($N \gg \hat{N}, M \gg \hat{M}$), overcoming the neural network's fixed structure constraint. Fig. 6 demonstrates this extension with $N = 10$ herders managing $M = 100$ targets using a policy originally trained for $\hat{N} = 2$, $\hat{M} = 5$ agents. Results show all targets ultimately reach the goal region, though some containment challenges
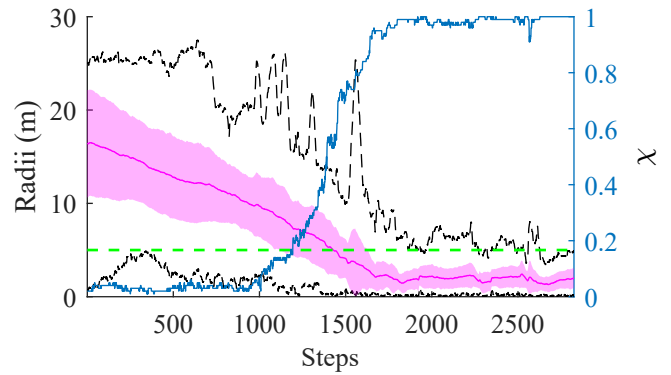

Fig. 6: Example of the policy behavior in a $N = 10, M = 100$ setting, where the previously trained $\hat{N} = 2, \hat{M} = 5$ policy is extended using topological sensing for the herders. The figure depicts the evolution of the targets' radii, showing the mean (magenta line), standard deviation (magenta shading), minimum and maximum values (black dashed lines) and the fraction of captured targets χ (blue line), alongside the goal region radius ρ_G (green dashed line).

exist due to limited sensing, similar to the herdability issues introduced and studied in [1]. Future work will focus on more robust, scalable approaches to address these remaining limitations.

VI. CONCLUSION

In this work, we advance decentralized learning-based shepherding control by leveraging Proximal Policy Optimization to handle continuous action spaces. Each herder operates independently without explicit communication, effectively solving the shepherding task in a fully model-free manner. Our solution demonstrates scalability to larger scenarios while maintaining a restricted observation space. The hierarchical architecture highlights the potential of advanced reinforcement learning in multi-agent control systems. However, the theoretical gap between our hierarchical solution and the optimal control strategy remains an open question. Future work will focus on: (1) extending to truly large-scale systems; (2) exploring heterogeneous agents with varying dynamics; (3) investigating adversarial target behaviors; (4) developing theoretical bounds on the optimality gap; and (5) validating with physical robots. Additionally, integrating limited communication protocols could potentially enhance coordination while maintaining computational efficiency.

APPENDIX

The parameters used for the numerical simulations described in Sections IV-V and the hyperparameters of the PPO and MAPPO algorithms are reported in Tables A1 and A2, respectively. Vidoes are available at https://github.com/SINCROgroup/Hierarchical-Policy-Gradient-RL-for-Multi-Agent-Shepherding-Control-of-Non-Cohesive-Targets.

TABLE A1: Simulation parameters for the shepherding model derived from [1]. The first-order model is integrated via Euler-Maruyama method with a sampling time $\Delta t = 0.05$ s.

Parameter	Value	Parameter	Value
ρ_G	5	D	0.5
R	25	λ	2.5
v_H	8	k^T	3

TABLE A2: Hyperparameters and reward gains of the PPO/MAPPO training algorithms (values for MAPPO are indicated in parentheses only when different from the PPO ones). The parameters names refer to the nomenclature found in [19].

Hyperparameter	Value	Gain	Value
Adam stepsize	5e-4	k_1	5e-2
Discount	0.98	k_2	1e-1
GAE parameter	0.95	k_3	1.5e-2
Clipping parameter	0.2	k_4	1e-2
VF coeff.	0.5		
Entropy coeff.	0.1 (0)		
Number of epochs	10		
Horizon	4096 (32)		
Minibatch size	128 (1024)		
Number of actors	8 (32)		

REFERENCES

[1] A. Lama and M. di Bernardo, "Shepherding and herdability in complex multiagent systems," *Phys. Rev. Res.*, vol. 6, p. L032012, 2024.

[2] A. Lama, M. di Bernardo, and S. H. L. Klapp, "An interpretable continuum framework for decision-making: nonreciprocal field theory of the shepherding control problem," *arXiv:2503.01112*, 2025.

[3] R. A. Licitra, Z. I. Bell, E. A. Doucette, and W. E. Dixon, "Single agent indirect herding of multiple targets: A switched adaptive control approach," *IEEE Control Systems Letters*, vol. 2, no. 1, pp. 127–132, 2018.

[4] A. A. Paranjape, S.-J. Chung, K. Kim, and D. H. Shim, "Robotic Herding of a Flock of Birds Using an Unmanned Aerial Vehicle," *IEEE Transactions on Robotics*, vol. 34, no. 4, pp. 901–915, 2018.

[5] J. Liu, H. Singh, S. Elsayed, R. Hunjet, and H. Abbass, "Planning-assisted context-sensitive autonomous shepherding of dispersed robotic swarms in obstacle-cluttered environments," *SSRN Electronic Journal*, 2023.

[6] V. S. Chipade, V. S. A. Marella, and D. Panagou, "Aerial Swarm Defense by StringNet Herding: Theory and Experiments," *Frontiers in Robotics and AI*, vol. 8, p. 640446, 2021.

[7] D. Ko and E. Zuazua, "Asymptotic behavior and control of a "guidance by repulsion" model," *Mathematical Models and Methods in Applied Sciences*, vol. 30, no. 04, pp. 765–804, 2020.

[8] D. Strömbom, R. Mann, A. Wilson, S. Hailes, A. Morton, D. Sumpter, and A. King, "Solving the shepherding problem: Heuristics for herding autonomous, interacting agents," *Journal of The Royal Society Interface*, vol. 11, 2014.

[9] J.-M. Lien, O. Bayazit, R. Sowell, S. Rodriguez, and N. Amato, "Shepherding behaviors," in *IEEE International Conference on Robotics and Automation*, vol. 4, pp. 4159–4164, 2004.

[10] J. Zhi and J.-M. Lien, "Learning to herd agents amongst obstacles: Training robust shepherding behaviors using deep reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 4163–4168, 2021.

[11] H. T. Nguyen, T. D. Nguyen, V. P. Tran, M. Garratt, K. Kasmarik, S. Anavatti, M. Barlow, and H. A. Abbass, "Continuous deep hierarchical reinforcement learning for ground-air swarm shepherding," *arXiv:2004.11543*, 2020.

[12] A. Hussein, E. Petraki, S. Elsawah, and H. A. Abbass, "Autonomous swarm shepherding using curriculum-based reinforcement learning," in *International Conference on Autonomous Agents and Multiagent Systems*, pp. 633–641, 2022.

[13] J. K. Gupta, M. Egorov, and M. Kochenderfer, "Cooperative multi-agent control using deep reinforcement learning," in *Autonomous Agents and Multiagent Systems*, pp. 66–83, 2017.

[14] G. Wang, J. Peng, C. Guan, J. Chen, and B. Guo, "Multi-drone collaborative shepherding through multi-task reinforcement learning," *IEEE Robotics and Automation Letters*, vol. 9, no. 11, pp. 10311–10318, 2024.

[15] Y. Hasan, J. E. G. Baxter, C. A. Salcedo, E. Delgado, and L. Tapia, "Flock navigation by coordinated shepherds via reinforcement learning," in *Algorithmic Foundations of Robotics XV*, pp. 454–469, 2022.

[16] F. Auletta, D. Fiore, M. J. Richardson, and M. di Bernardo, "Herding stochastic autonomous agents via local control rules and online target selection strategies," *Autonomous Robots*, vol. 46, no. 3, pp. 469–481, 2022.

[17] F. de Lellis, F. Auletta, G. Russo, P. de Lellis, and M. di Bernardo, "An application of control- tutored reinforcement learning to the herding problem," in *17th International Workshop on Cellular Nanoscale Networks and their Applications*, pp. 1–4, 2021.

[18] I. Napolitano, A. Lama, F. De Lellis, and M. di Bernardo, "Emergent cooperative strategies for multi-agent shepherding via reinforcement learning," *arXiv:2411.05454*, 2024.

[19] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv:1707.06347*, 2017.

[20] G. Albi, M. Bongini, E. Cristiani, and D. Kalise, "Invisible control of self-organizing agents leaving unknown environments," *SIAM Journal on Applied Mathematics*, vol. 76, no. 4, pp. 1683–1710, 2016.

[21] P. Dolai, A. Simha, and S. Mishra, "Phase separation in binary mixtures of active and passive particles," *Soft Matter*, vol. 14, no. 29, pp. 6137–6145, 2018.

[22] A. Pierson and M. Schwager, "Controlling Noncooperative Herds with Robotic Herders," *IEEE Transactions on Robotics*, vol. 34, no. 2, pp. 517–525, 2018.

[23] A. Li, M. Ogura, and N. Wakamiya, "Communication-free shepherding navigation with multiple steering agents," *Frontiers in Control Engineering*, vol. 4, p. 989232, 2023.

[24] C. F. Nino, O. S. Patil, J. N. Philor, Z. I. Bell, and W. E. Dixon, "Deep adaptive indirect herding of multiple target agents with unknown interaction dynamics," in *62nd IEEE Conference on Decision and Control*, pp. 2509–2514, 2023.

[25] C. Yu, A. Velu, E. Vinitsky, J. Gao, Y. Wang, A. Bayen, and Y. Wu, "The surprising effectiveness of ppo in cooperative multi-agent games," *Advances in neural information processing systems*, vol. 35, pp. 24611–24624, 2022.

[26] N. Heess, D. TB, S. Sriram, J. Lemmon, J. Merel, G. Wayne, Y. Tassa, T. Erez, Z. Wang, S. M. A. Eslami, M. Riedmiller, and D. Silver, "Emergence of Locomotion Behaviours in Rich Environments," *arXiv:1707.02286*, 2017.

[27] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *32nd International Conference on Machine Learning*, pp. 1889–1897, 2015.

[28] M. Andrychowicz, A. Raichuk, P. Stańczyk, M. Orsini, S. Girgin, R. Marinier, L. Hussenot, M. Geist, O. Pietquin, M. Michalski, S. Gelly, and O. Bachem, "What Matters for On-Policy Deep Actor-Critic Methods? A Large-Scale Study," in *International Conference on Learning Representations*, 2020.

[29] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic, "Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study," *Proceedings of the National Academy of Sciences*, vol. 105, no. 4, pp. 1232–1237, 2008.

[30] Anthropic, "Claude (version 3.7 sonnet)." Anthropic, 2025. AI assistant used for English language revision.