



Herding stochastic autonomous agents via local control rules and online target selection strategies

Fabrizia Auletta^{1,2} · Davide Fiore³ · Michael J. Richardson^{2,4} · Mario di Bernardo⁵

Received: 4 May 2021 / Accepted: 20 December 2021 / Published online: 16 February 2022
© The Author(s) 2022

Abstract

We propose a simple yet effective set of local control rules to make a small group of “herder agents” collect and contain in a desired region a large ensemble of non-cooperative, non-flocking stochastic “target agents” in the plane. We investigate the robustness of the proposed strategies to variations of the number of target agents and the strength of the repulsive force they feel when in proximity of the herders. The effectiveness of the proposed approach is confirmed in both simulations in ROS and experiments on real robots.

Keywords Agent-based systems · Biologically-inspired agents · Autonomous agents · Multi-robot systems

The authors wish to acknowledge support from the Macquarie Cotutelle (Industrial and International Leverage Fund) Award from University of Bristol and International Macquarie University Research Excellence Scholarship Scheme from Macquarie University for supporting Fabrizia Auletta’s work. This research was supported, in part, by Australian Research Council Future Fellowship (FT180100447) awarded to Michael Richardson, in part with the economic support of MIUR (Italian Ministry of University and Research) performing the activities of the Project ARS01_00861 “Integrated collaborative systems for smart factory—ICOSAF”, and in part by the University of Naples Federico II—“Finanziamento della Ricerca di Ateneo (FRA) Linea B” through the research grant “BIOMASS”.

✉ Mario di Bernardo
mario.dibernardo@unina.it

Fabrizia Auletta
fabrizia.auletta@bristol.ac.uk

Davide Fiore
davide.fiore@unina.it

Michael J. Richardson
michael.j.richardson@mq.edu.au

¹ Department of Engineering Mathematics, University of Bristol, University Walk, Bristol BS8 1TR, UK

² School of Psychological Sciences, Faculty of Medicine, Health and Human Sciences, Macquarie University, Sydney, NSW 2109, Australia

³ Department of Mathematics and Applications “R. Caccioppoli”, University of Naples Federico II, Via Cintia, Monte S. Angelo, 80126 Naples, Italy

1 Introduction

Exploration and rescue, evacuation from dangers, surveillance and crowd control are all examples of *multi-agent herding problems* in which two kinds of agents interact.

In these problems, a set of “active” agents (the herders) need to drive a set of “passive” agents (the herd) towards a desired goal region and confine them therein (Long et al. 2020; Nolfi 2002).

In most cases, repulsive forces exerted by the herders on the herd are exploited to drive the movements of the passive agents that need to be corralled and, at times, cooperation among the herders (such as attractive forces between them) are used to enhance the herding performance.

Notable herding solutions are those proposed in Vaughan et al. (2000), Lien et al. (2004), Strombom et al. (2014), Paranjape et al. (2018), Licitra et al. (2019) for single herders and in Lien et al. (2005), Haque et al. (2009), Lee and Kim (2017), Pierson and Schwager (2018), Nalepka et al. (2017b), Montijano et al. (2013), Varava et al. (2017), Song et al. (2021), Chipade and Panagou (2019), Sebastián and Montijano (2021) for multiple herders.

⁴ The Center for Elite Performance, Expertise and Training, Macquarie University, Sydney, NSW 2109, Australia

⁵ Department of Electrical Engineering and Information Technology, University of Naples Federico II, Via Claudio 21, 80125 Naples, Italy

One of the problems to be addressed in designing control strategies to drive herder agents is endowing them with the ability of deciding at any given time what passive agent a herder should target first when multiple herders and targets are present.

For the sake of comparison with our approach, we now briefly review the most relevant research from the literature addressing multi-agent herding, where multiple herders are required to collect and drive a group of passive agents towards a desired goal region.

1.1 Related work

One of the earliest solutions to the herding problem was proposed by Lien et al. (2004, 2005). The trajectories followed by passive and herder agents were generated using global rule-based roadmaps—abstract representations of the walkable paths given as a directed graph (Wilmarth et al. 1999). Numerical simulations showed that multiple herders were successful in coping with increasing sizes of the herd. Nevertheless, herders' performance worsened as the flocking tendency of passive agents decreased.

Multi-agent herding scenarios were also considered in Haque et al. (2009, 2011). Here the authors addressed the problem of controlling a group of herders so as to entrap a group of passive agents in a region from which they could not escape. To solve this problem, each herder was pre-assigned some region of influence. Targets' motion was then only influenced by a specific herder if they happened to be within its region of influence; targets travelling otherwise at constant speed and with a heading aligned to that of their neighbouring agents. The velocities of the herders were regulated according to that of the other passive agents with which they interacted.

Other multi-agent herding scenario, where many herders are required to collect and patrol a group of passive agents, were also proposed in Lee and Kim (2017). Inspired by the limited visual field of real sheepdogs and the absence of centralised coordination among them, the latter work proposed a herding algorithm based entirely on local control rules. The dynamics of both herders and passive agents were modelled as the linear combination of potential field-like forces within a sensing area. In addition to this basic dynamics, passive agents were also subject to a repulsive force from the herders. Herders were controlled by an appropriate input selected as a function of their distance from the nearest passive agent and their distance from a desired goal. The result of the proposed shepherding behaviour was the emergence of an arc formation among the herders [a similar formation was instead hard coded in the algorithm presented earlier in Lien et al. (2005)]. Numerical simulations showed the effectiveness of the approach under the assumption that passive agents tend to flock together. In this case, herders could indeed collect and

herd multiple sub-flocks without any explicit coordination rule.

In Robotics, feedback control strategies have been recently presented to solve multi-agent herding problems and guarantee convergence of the overall system. For instance, in Pierson and Schwager (2018) the case of multiple herder agents regulating the mean position of a group of flocking passive agents was investigated. An arc-based strategy was proposed for the herders to surround and drive the targets towards a desired goal region. The proposed control law and its convergence properties were explored by modelling the whole herd as a single unicycle controlled by means of a point-offset technique. Montijano et al. (2013) proposed a herding strategy based on elliptical orbits to entrap a passive agent whose position is uncertain. Varava et al. (2017); Song et al. (2021) developed a “herding by caging” solution, based on geometric considerations and motion planning techniques to arrange the herder agents around the flock. A similar formation was presented in Chipade and Panagou (2019), and further developed in Chipade et al. (2021), to let herders identify clusters of flocking adversarial agents, dynamically encircle and drive them to a safe zone. Recently, Sebastián and Montijano (2021) developed analytical and numerical control design procedures to compute suitable herding actions to herd evading agents to a desired position, even when the nonlinearities in the evaders' dynamics yield implicit equations.

A different approach was used in Cognitive Science (Nalepka et al. 2015, 2017a, b, 2019), where a model of the herding agent was derived from experimental observations of how two human players herd a group of randomly moving agents in a virtual reality setting. It was observed that, at the beginning of the task, all pairs of human players adopt a *search and recovery* strategy; players individually chasing the farthest passive agent in the half of the game field assigned to them, driving it inside the desired containment region. Once all agents are gathered inside the goal region, most pairs of human herders were observed to switch to an entirely different containment strategy, based on exhibiting an oscillatory movement along an arc around the goal region creating effectively a “repulsive wall” for the passive agents keeping them therein (Nalepka et al. 2017a). To reproduce this behaviour in artificial agents, a nonlinear model was proposed by Nalepka et al. (2019) where the switch from search and recovery to the oscillatory containment strategy is induced by a Hopf bifurcation triggered by a change in the distance of the herd agents from the goal region.

With regard to a *single* herder agent gathering one-by-one a group of passive agents, recent work by Licitra et al. (2017) employed a backstepping control strategy for the single herder to chase one target at a time, with the herder switching among different targets and succeeding in collecting them within a goal region of interest. This idea was further

developed in Licitra et al. (2018, 2019) where other control strategies and further uncertainties in the herd's dynamics were investigated.

1.2 Contributions of this paper

In this paper, we consider the case of a small group of herders chasing a much larger group of passive agents whose dynamics, as often happens with natural agents such as fish, birds or bacteria, is stochastic and driven by a random Brownian noise. However, contrary to what is usually done in the rest of the literature (Haque et al. 2009; Lien et al. 2004; Pierson and Schwager 2018; Lee and Kim 2017; Chipade and Panagou 2019), we do not consider the presence of any flocking behaviour between passive agents, making the problem more complicated to solve as each target needs to be tracked and collected independently from the others.

To solve the problem, we present a simple, yet effective, dynamic herding strategy consisting of local feedback control laws for the herder agents and a set of target selection rules that drive how herders make decentralised decisions on what targets to follow. A herder's action is based on the global knowledge of the environment and of the positions of all other agents. With respect to other solutions in the literature (Lien et al. 2004; Pierson and Schwager 2018; Chipade et al. 2021; Song et al. 2021), our approach does not involve the use of *ad hoc* formation control strategies to force the herders surround the herd, but we rather enforce cooperation between herders by dynamically dividing the plane among them by means of simple yet effective and robust rules that can be easily implemented in real robots.

We then numerically analyse how robust these strategies are to parameter perturbations, uncertainties and unmodeled disturbances in passive agent dynamics. Moreover, we assess how different choices of the target selection rule affect the overall effectiveness of the methodology we propose. Finally, we test the effectiveness of the proposed strategies to solve the herding problem firstly in simulations in ROS and then in experiments on real robots conducted via the online Robotarium platform (Pickem et al. 2017; Wilson et al. 2020).

2 The herding problem

We consider the problem of controlling $N_H \geq 2$ herder agents in order for them to drive a group of $N_T > N_H$ passive agents in the plane (\mathbb{R}^2) towards a goal region and contain them therein. We term \mathbf{y}_j the position in Cartesian coordinates of the j -th herder in the plane and \mathbf{x}_i that of the i -th passive agent. We denote as (r_j, θ_j) and (ρ_i, ϕ_i) their respective positions in polar coordinates as shown in Fig. 1. We assume the goal of the herders is to drive the passive agents towards a circular *containment region* \mathcal{G} , of radius r^*




Fig. 1 Illustration of the spatial arrangement in the herding problem. The herder agent \mathbf{y}_j (yellow square), with polar coordinates (r_j, θ_j) , must relocate the target agent \mathbf{x}_i (green ball), with polar coordinates (ρ_i, ϕ_i) , in the containment region \mathcal{G} (solid red circle) of center \mathbf{x}^* and radius r^* . The buffer region \mathcal{B} , of width Δr^* , is depicted as a dashed red circle (Color figure online)

centred at \mathbf{x}^* . Without loss of generality, we set \mathbf{x}^* to be the origin of \mathbb{R}^2 .

Assuming the herders have their own trivial dynamics in the plane, the *herding problem* can be formulated as that of designing the control action u governing the dynamics of the herders assumed to be given by

$$m \ddot{\mathbf{y}}_j = u(t, \mathbf{x}_1, \dots, \mathbf{x}_{N_T}, \mathbf{y}_1, \dots, \mathbf{y}_{N_H}), \quad (1)$$

where m denotes the mass of the herders taken to be unitary, so that the herders can influence the dynamics of the passive agents (whose dynamics will be specified in the next section) and guarantee that

$$\|\mathbf{x}_i(t) - \mathbf{x}^*\| \leq r^*, \quad \forall i, \forall t \geq t_g, \quad (2)$$

where $\|\cdot\|$ denotes the Euclidean norm; that is, all passive agents are contained, after some finite *gathering time* t_g , in the desired region \mathcal{G} . A herding trial is said to be *successful* in the time interval $[0, T]$ if condition (2) holds for some $t_g \in [0, T]$. We assume an annular *safety region* \mathcal{B} of width Δr^* exists surrounding the goal region that the herders leave between themselves and the region where targets are contained.

In what follows, we will assume that (i) herder and passive agents can move freely in \mathbb{R}^2 ; (ii) herder agents have global knowledge of the environment and of the positions of the other agents therein.

3 Target dynamics

Taking inspiration from Nalepka et al. (2017b), we assume that, when interacting with the herders, passive agents are repelled from them and move away in the opposite direction, while in the absence of any external interaction, they randomly diffuse in the plane. Specifically, we assume passive agents move according to the following stochastic dynamics

$$d\mathbf{x}_i(t) = \mathbf{V}_{r,i}(t)dt + \alpha_b d\mathbf{W}_i(t), \tag{3}$$

where $\mathbf{V}_{r,i}$ describes the repulsion exerted by all the herders on the i -th passive agent, $\mathbf{W}_i = [W_{i,1}, W_{i,2}]^\top$ is a 2-dimensional standard Wiener process and $\alpha_b > 0$ is a constant. We suppose the distance travelled by the passive agents depends on how close the herder agents are and model this effect by considering a potential field centred on the j -th herder given by $v_{i,j} = 1/(\|\mathbf{x}_i - \mathbf{y}_j\|)$, exerting on the passive agents an action proportional to its gradient (Pierson and Schwager 2018). Specifically, the dynamics of the i -th passive agent is influenced by all the herders through the reaction term

$$\mathbf{V}_{r,i}(t) = \alpha_r \sum_{j=1}^{N_H} \frac{\partial v_{i,j}}{\partial \mathbf{x}_i} = -\alpha_r \sum_{j=1}^{N_H} \frac{\mathbf{x}_i(t) - \mathbf{y}_j(t)}{\|\mathbf{x}_i(t) - \mathbf{y}_j(t)\|^3}, \tag{4}$$

where $\alpha_r > 0$ is a constant. Uncertainties on the repulsive reaction term (4) can be seen as being captured by the additional noisy term in (3).

Notice that the velocity of all passive agents is completely determined by (3)–(4) and we do not assume any upper bound on its maximum value. The position of the i -th passive agent when it is targeted by the j -th herder will be denoted as $\tilde{\mathbf{x}}_{i,j}$ or in polar coordinates as $(\tilde{\rho}_{i,j}, \tilde{\phi}_{i,j})$.

4 Herder dynamics and control rules

Our solution to the herding problem consists of two layered strategies; (i) a local control law to drive the motion of the herder towards the target it selected, and to push it inside the goal region and (ii) a target selection strategy through which herders decide what target they need to chase. When the herd are all gathered, the herders switch back to an idling condition by keeping themselves within the safety region surrounding the goal region.

4.1 Local control strategy

For the sake of comparison with the strategy presented in Nalepka et al. (2017b), Nalepka et al. (2019), we derive

in polar coordinates the control law we propose to drive each herder. Albeit not resulting in the shortest possible path travelled by the herders, the controller expressed in polar coordinates ensures circumnavigation of the goal region, avoiding passive agents already contained therein from being scattered around. Specifically, the control input to the j -th herder dynamics (1) is defined as $\mathbf{u}_j = u_{r,j} \hat{r}_j + u_{\theta,j} \hat{\theta}_j$, where $\hat{\theta}_j = \hat{r}_j^\perp$ are unit vectors and $\hat{r}_j = [\cos \theta_j, \sin \theta_j]^\top$, and its components are chosen as

$$u_{r,j}(t) = -b_r \dot{r}_j(t) - \mathcal{R}(\tilde{\mathbf{x}}_{i,j}, t), \tag{5}$$

$$u_{\theta,j}(t) = -b_\theta \dot{\theta}_j(t) - \mathcal{T}(\tilde{\mathbf{x}}_{i,j}, t), \tag{6}$$

with $b_r, b_\theta > 0$, and where the feedback terms $\mathcal{R}(\tilde{\mathbf{x}}_{i,j}, t)$ and $\mathcal{T}(\tilde{\mathbf{x}}_{i,j}, t)$ are elastic forces that drive the herder towards the target i and push it towards the containment region \mathcal{G} . Such forces are chosen as

$$\mathcal{R}(\tilde{\mathbf{x}}_{i,j}, t) = \epsilon_r \left[r_j(t) - \xi_j(t) (\tilde{\rho}_{i,j}(t) + \Delta r^*) - (1 - \xi_j(t)) (r^* + \Delta r^*) \right], \tag{7}$$

$$\mathcal{T}(\tilde{\mathbf{x}}_{i,j}, t) = \epsilon_\theta \left[\theta_j(t) - \xi_j(t) \tilde{\phi}_{i,j}(t) - (1 - \xi_j(t)) \psi(t) \right]. \tag{8}$$

with $\epsilon_r, \epsilon_\theta > 0$, and where ξ_j regulates the switching policy between collecting and idling behaviours. That is, $\xi_j = 1$, if $\tilde{\rho}_{i,j} \geq r^*$, and $\xi_j = 0$, if $\tilde{\rho}_{i,j} < r^*$, so that the herder is attracted to the position of the i -th target $\tilde{\mathbf{x}}_{i,j}$ (plus a radial offset Δr^*) when the current target is outside the containment region ($\xi_j = 1$) or close to the boundary of the buffer region at the idling position $(r^* + \Delta r^*, \psi)$, in polar coordinates, otherwise ($\xi_j = 0$). The value of the *idling angle* ψ depends on the specific choice of the target selection strategy employed, which are discussed next. Note that the control laws (5)–(6) are much simpler than those presented by Nalepka et al. (2017b) as they do not contain any higher order nonlinear term nor are complemented by parameter adaptation rules (see Nalepka et al. 2017b for further details). Moreover, as for the passive agents, we do not assume any upper bound on the maximum velocity of the herders.

4.2 Target selection strategies

In the case of a single herder chasing multiple agents, the most common strategy in the literature is for it to select the target as either the farthest passive agent from the goal region, or the centre of mass of the flocking herd (Vaughan et al. 2000; Strombom et al. 2014; Licitra et al. 2017). When two or more herders are involved, the problem is usually solved using a formation control approach, letting the herders surround the herd and then drive them towards the goal region (Pierson




Fig. 2 Graphical representation of the target selection strategies. Herders are depicted as yellow squares, passive agents as green balls. The colours in which the game field is divided correspond to regions assigned to different herders. Herder y_j is currently chasing target agent $\tilde{x}_{i,j}$, while passive agent x_i is not chased by any herder (Color figure online)

and Schwager 2018; Lien et al. 2004). Rather than using formation control techniques or solving off-line or on-line optimisation problems as in dynamic target assignment problems (e.g., Bürger et al. 2011), here we present a set of simple, yet effective, target selection strategies that exploit the spatial distribution of the herders allowing them to cooperatively select their targets without requiring any computationally expensive optimisation problem to be solved on-line.

We present four different herding strategies, starting from the simplest case where herders globally look for the target farthest from the goal region. A graphical illustration of the four strategies is reported in Fig. 2 for $N_H = 3$ herders. *Global search strategy (no partitioning)*. Each herder selects the farthest passive agent from the containment region which is not currently targeted by any other herder (Fig. 2a). Being the simplest, we present this strategy for the sake of comparison only and not for being implemented on real robots.

Static arena partitioning At the beginning of the trial and for all of its duration, the plane is partitioned in N_H circular sectors of width equal to $2\pi/N_H$ rad centred at x^* . Each herder is then assigned one sector to patrol and selects the passive agent therein that is farthest from \mathcal{G} (Fig. 2b). Note

that this is the same herding strategy used in Nalepka et al. (2017b) for $N_H = 2$ herders.

Dynamic leader-follower (LF) target selection strategy. At the beginning of the trial, herders are labelled from 1 to N_H in anticlockwise order starting from a randomly selected herder which is assigned the *leader* role. The plane is then partitioned dynamically in different regions as follows. The leader starts by selecting the farthest passive agent from \mathcal{G} whose angular position $\phi_{i,1}$ is such that

$$\tilde{\phi}_{i,1} \in \left(\theta_1 - \frac{1}{2} \frac{2\pi}{N_H}, \theta_1 + \frac{1}{2} \frac{2\pi}{N_H} \right],$$

where θ_1 is the angular position of the leader at time t . Then, all the other *follower* herders ($j = 2, \dots, N_H$), in ascending order, select their targets as the passive agent farthest from \mathcal{G} such that

$$\tilde{\phi}_{i,j} \in \left(\theta_1 - \frac{1}{2} \frac{2\pi}{N_H} + \zeta_j, \theta_1 + \frac{1}{2} \frac{2\pi}{N_H} + \zeta_j \right],$$

with $\zeta_j = 2\pi(j - 1)/N_H$. As the leader chases the selected target and moves in the plane, the partition described above changes dynamically so that a different circular sector with *constant* angular width $2\pi/N_H$ rad is assigned to each follower at any time instant. In Fig. 2c the case is depicted for $N_H = 3$ in which the sector $(\theta_1 - \frac{\pi}{3}, \theta_1 + \frac{\pi}{3})$ is assigned to the leader herder (assumed to be y_1) while the rest of the plane is assigned equally to the other two herders.

Dynamic peer-to-peer (P2P) target selection strategy. At the beginning of the trial herders are labelled from 1 to N_H as in the previous strategy. Denoting as ζ_j^+ the angular difference between the positions of herder j and herder $(j + 1) \bmod N_H$ at time t , and as ζ_j^- that between herder j and herder $(j + N_H - 1) \bmod N_H$ at time t , then herder j selects the farthest passive agent from \mathcal{G} whose angular position is such that

$$\tilde{\phi}_{i,j} \in \left(\theta_j - \frac{\zeta_j^-}{2}, \theta_j + \frac{\zeta_j^+}{2} \right].$$

Unlike the previous case, now the width of the circular sector assigned to each herder is also dynamically changing as it depends on the relative angular positions of the herders in the plane.

The *idling angle* ψ in (8) is set equal to the angular position $\tilde{\phi}_{i,j}$ of the last contained target for the *global search strategy*, otherwise it is set equal to the angular position corresponding to the half of the angular sector assigned at each time to the herder. In so doing, the herder is made to rest at a point in which its angular distance is minimised from any passive agent escaping the containment region into the assigned angular sector.

A crucial difference between the herding strategies presented above is the nature (local vs global) and amount of information that herders must possess to select their next target. Specifically, when the *global search strategy* is used, every herder needs to know the position \mathbf{x}_i of every passive agent in the plane, not currently targeted by other herders. In the case of the *static arena partitioning* instead a herder needs to know its assigned (constant) circular sector together with the position \mathbf{x}_i of every passive agent in the sector.

For the dynamic target selection strategies, less information is generally required. Indeed, in the *dynamic leader-follower strategy* the herders, knowing N_H , can either self-select the sector assigned to themselves (if they act as leader) or self-determine their respective sector by knowing the position of the leader \mathbf{y}_1 . Similarly in the *dynamic peer-to-peer strategy* herders can self-select their sectors by using the angles ζ_j^+ and ζ_j^- .

Note that in the event of perfect radial alignment of the herder and its target, the herder might push the target away, rather than towards the goal region (Fig. 3).

Although this condition is very unlikely to persist due to the random motion of the passive agents, this problem can be avoided by extending the herder dynamics in (1) by a circumnavigation force $u_j^\perp(t)$. This force is orthogonal to the vector $\Delta\mathbf{x}_{ij} = \mathbf{x}_i - \mathbf{y}_j$, and its amplitude depends on the angle χ_{ij} between $\Delta\mathbf{x}_{ij}$ and \mathbf{y}_j , such that it is maximal when the two vectors are parallel ($\chi_{ij} = 0$) and zero when they are anti-parallel ($\chi_{ij} = \pi$). Specifically, it is defined as:

$$u_j^\perp(t) = \bar{U} \cdot v(t) \cdot \cos^2\left(\frac{\chi_{ij}}{2}\right) \frac{\Delta\mathbf{x}_{ij}^\perp}{\|\Delta\mathbf{x}_{ij}\|}, \tag{9}$$

where $\bar{U} > 0$ is the maximum amplitude, and the value of $v \in \{-1, 1\}$ depends on which halves of the assigned sector the herder is currently in to guarantee that the target agent is always pushed toward the interior of the sector.

5 Numerical validation

The herding performance of the proposed control strategies has been evaluated through a set of numerical experiments aimed at (i) assessing their effectiveness in achieving the herding goal; (ii) comparing the use of different target selection strategies; (iii) studying the robustness of each strategy to parameter variations. The implementation and validation of the strategies in a more realistic robotic environment is reported in the next section where ROS simulations and experiments are included.




Fig. 3 Graphical representation of circumnavigation force $u_j^\perp(t)$ in the case that herder and its target are not aligned (a $\chi_{ij} \in (0, \pi)$) and in the case the herder is perfectly aligned behind (b $\chi_{ij} = \pi$) or ahead (c $\chi_{ij} = 0$) of the passive agent w.r.t. the containment region \mathcal{G}

5.1 Performance metrics

We defined the following metrics (see Appendix A for their definitions) to evaluate the performance of different strategies. Specifically, for each of the proposed strategies we computed the (i) gathering time t_g , (ii) the average length d_g of the path travelled by the herders until all targets are contained, (iii) the average total length d_{tot} of the path travelled by herders during all the herding trial, (iv) the mean distance D_T between the herd’s centre of mass and the centre of the containment region, and (v) the herd agents’ spread $S\%$.

Note that *lower* values of t_g correspond to *better* herding performance; herders taking a shorter time to gather all the passive agents in the goal region. Also, *lower* values of D_T and $S\%$ correspond to a *tighter* containment of the passive agents in the goal region while *lower* values of d_g and d_{tot} correspond to a *more efficient* herding capability of the herders during the gathering and containment of the herd.

5.2 Performance analysis

We carried out 50 simulation trials with $N_T = 7$ passive agents and either $N_H = 2$ or $N_H = 3$ herders, starting from random initial conditions. All simulation trials were found

Table 1 Average performance and standard deviation over 50 successful trials of different herding strategies for $N_T = 7$ passive agents

	Global	Static	LF	P2P
$N_H = 2$				
t_g [a.u.]	8.5 ± 1.5	15.2 ± 9.6	15.3 ± 9.2	13.3 ± 6.2
d_g [a.u.]	139 ± 34	102 ± 42	92 ± 49	143 ± 53
d_{tot} [a.u.]	841 ± 27	493 ± 51	423 ± 29	418 ± 56
D_T [a.u.]	1.3 ± 0.1	1.4 ± 0.6	1.5 ± 0.6	1.3 ± 0.5
$S\%$ [%]	0.1 ± 0.05	0.2 ± 0.1	0.2 ± 0.2	0.2 ± 0.1
$N_H = 3$				
t_g [a.u.]	5.8 ± 0.9	11.2 ± 7.9	11.2 ± 8.7	8.3 ± 5.6
d_g [a.u.]	88 ± 21	107 ± 98	84 ± 74	58 ± 43
d_{tot} [a.u.]	1242 ± 26	757 ± 67	885 ± 56	950 ± 79
D_T [a.u.]	0.6 ± 0.1	0.9 ± 0.5	0.8 ± 0.3	0.7 ± 0.1
$S\%$ [%]	0.1 ± 0.03	0.2 ± 0.2	0.2 ± 0.2	0.2 ± 0.3

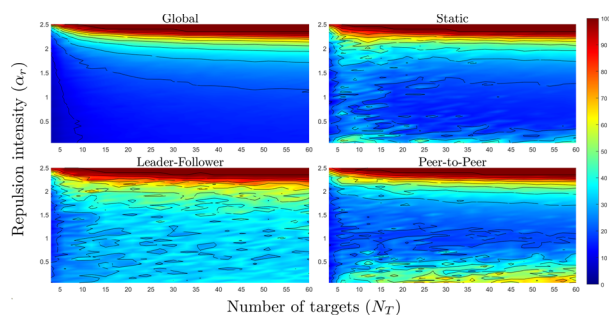
to be successful, that is, such that condition (2) is verified. (All simulation parameters and the description of simulation setup adopted here are reported in “Appendix B”).

The results of our numerical investigation are reported in Table 1. As expected, when herders search globally for agents to chase, their average total path, d_{tot} , is notably larger than when dynamic target selection strategies are used, pointing out that this strategy is going to be the least efficient when implemented and also requiring complete information about the agents. Therefore, in what follows we will discuss this strategy only for the sake of completeness and not for the purpose of its implementation.


As regards the aggregation of the herd in terms of D_T and $S\%$, all other strategies presented comparable results in terms of both mean and standard deviation. Dynamic strategies showed better gathering performance (t_g and d_g) than the static arena partitioning. Therefore, we find that in general higher level of cooperation between herders and a more efficient coverage of the plane, as those guaranteed by dynamic strategies, yield an overall better herding performance which is more suitable for realistic implementations in robots or virtual agents that are bound to move at limited speed.

5.3 Robustness analysis

Next, we analysed the robustness of the proposed herding strategies to variations of the herd size and of the magnitude of the repulsive reaction to the herders exhibited by the passive agents (Fig. 4). Specifically, we varied N_T between 3 and 60 and the repulsion parameter α_r in (4) between 0.05 and 2.5, while keeping $N_H = 2$; we found that all strategies succeed in herding up to 60 agents in a large region of parameter values (see the blue areas in Fig. 4a). The global strategy, where herders patrol the entire plane, is found as expected to



(a) Gathering time t_g . Lower values correspond to faster herding.



(b) Total distance travelled d_{tot} . Lower values correspond to more efficient herding.

Fig. 4 Robustness analysis of the proposed herding strategies for two herders ($N_H = 2$) to variation of herd size N_T and repulsive reaction coefficient α_r . N_T was varied between 3 and 60 agents, with increments equal to 3, while α_r between 0.05 and 2.5, with increments equal to 0.05. For each pair (N_T, α_r) the corresponding metric was averaged over 15 simulation trials starting with random initial positions. The coloured plots were obtained by interpolation of the computed values (Color figure online)

be the least efficient in terms of total distance traveled by the herders (Fig. 4b); the dynamic peer-to-peer strategy offering the best compromise and robustness property in terms of containment performance (Fig. 4a) and efficiency (Fig. 4b).

To further validate these findings we carried out 50 simulation trials for the herding scenario in which $N_H = 3$ herders are required to herd $N_T = 60$ passive agents.

In this more challenging scenario not all the trials were found to be successful, as per (2), due to at least one passive agent escaping the containment region, and we averaged the resulting performance over the successful trials only (Table 2). Herders adopting the global and peer-to-peer strategies successfully herd all agents in over 50% of the trials. Moreover, herders globally searching for the target to chase spent on average slightly less time gathering the targets ($t_g = 12.96$) and achieved and maintained lower herd spread ($S\% = 0.48$). However, the path travelled to achieve the goal (d_{tot}) was significantly larger than when static or dynamic selection strategies were adopted.

Table 2 Average performance over successful trials of different herding strategies for $N_T = 60$ passive agents

	Global	Static	LF	P2P
$N_H = 3$				
Successful trials	49	11	7	26
t_g [a.u.]	12.9 ± 1.8	18.2 ± 4.2	16.1 ± 5.1	15.9 ± 6.6
d_g [a.u.]	211 ± 23	191 ± 43	138 ± 31	148 ± 66
d_{tot} [a.u.]	1226 ± 25	708 ± 28	742 ± 74	734 ± 97
D_T [a.u.]	6.5 ± 0.7	10.8 ± 4.1	8.2 ± 3.3	8.9 ± 6.8
$S_{\%}$ [%]	0.5 ± 0.05	0.9 ± 1.5	3.4 ± 5.8	1.1 ± 1.7

6 Experimental validation

To validate in more realistic robotic settings the strategies we propose, we complemented the numerical simulations presented in Sect. 5 with simulations in ROS¹ and experiments on real robots conducted on the online Robotarium platform (Pickem et al. 2017; Wilson et al. 2020).

6.1 ROS simulations

ROS² is an advanced framework for robot software development that provides tools to support the user during all the development cycle, from low-level control and communication to deployment on real robots. We used the Gazebo software package³ to test the designed control architecture on accurate 3D models of commercial robots to simulate their dynamics and physical interaction with the virtual environment.

We considered a scenario where $N_T = 3$ passive agents need to be herded by $N_H = 2$ robotic herders. All agents were chosen to be implemented as Pioneer 3-DX,⁴ a commercially available two-wheel two-motor differential drive robot whose detailed model is available in Gazebo (Fig. 5). The desired trajectories for the robots are generated by using Eqs. (3) and (5)–(6) for the passive and herder robots, respectively, which are used as reference signals for the on-board inner control loop to generate the required tangential and angular velocities (see “Appendix C” for further details).

Examples of ROS simulations are reported in Fig. 6 where all the target selection strategies that were tested (static arena partitioning, leader-follower, peer-to-peer) were found to be successful with herder robots being able to gather all the passive robots in the containment region. Figure 6 also shows that the angular position of the herders remain within the




Fig. 5 Overview of Gazebo-ROS application, with 3D model of the Pioneer 3-DX robot (a) and a landscape view of the simulated environment (b)

bounds defining the sector of the plane assigned to them for patrolling.

6.2 Robotarium experiments

Robotarium is a remotely accessible swarm robotics research platform equipped with GRITBot robots which allows rapid deployment and testing of custom control algorithms (Pickem et al. 2015; Wilson et al. 2020). To comply with the limited space of the arena ($3.2 \text{ m} \times 2 \text{ m}$) and safety protocols to avoid collisions between robots (robots’ diameter is 11 cm) implemented in the platform, we considered a scenario with $N_T = 4$ passive robots and $N_H = 2$ herder robots; a herding scenario that was also considered in Rigoli et al. (2020), Auletta et al. (2021) to study and model the selection strategies adopted by pairs of human-driven herder agents.

Herder parameters were selected as described in Appendix B, while the coefficient of diffusion and repulsion in the dynamics of passive agents (3) were scaled to $(\alpha_b, \alpha_r) = (0.001, 0.4)$ to comply with the physical constraints on the hardware of the GRITBots; having a max tangential speed of 20 cm/s and a max rotational speed of about 3.6 rad/s. The results of the experimental test are reported in Fig. 7. Both dynamic strategies were found to be effective in containing all the passive robots with a gathering time t_g less than 70 s, and with the peer-to-peer strategy guaranteeing slightly faster convergence (Fig. 7c) than the leader-follower strategy (Fig. 7b) over all the 5 trials that were

¹ Code available at <https://github.com/diBernardoGroup/HerdingProblem>.

² <https://www.ros.org>.

³ http://wiki.ros.org/gazebo_ros_pkgs.

⁴ Pioneer 3 - Operations Manual, available at https://www.inf.ufrgs.br/~prestes/Courses/Robotics/manual_pioneer.pdf (2020/08/11).

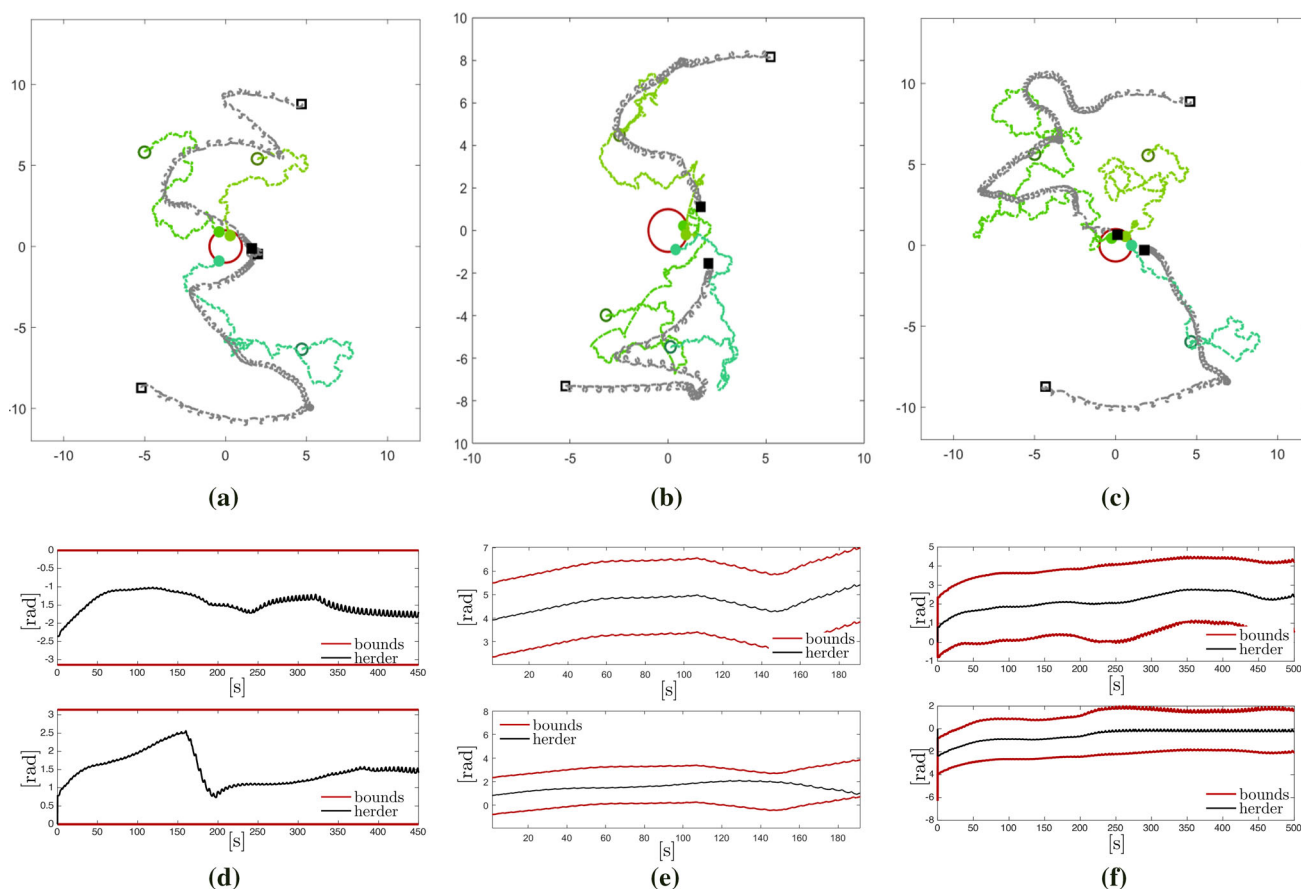


Fig. 6 ROS simulations. Top panels show the trajectories of passive agents (green lines) and herders (grey lines) adopting **a** static arena partitioning, **b** leader-follower and **c** peer-to-peer herding strategies simulated in the Gazebo environment. The containment region \mathcal{G} is depicted as a red circle. Black square marks denote the initial and the final (solid coloured) position of the herders. Green circle marks show

the initial and the final (solid coloured) position of the passive agents. Bottom panels show that all herders are able to collect the herd in less than 500 s by following the angular bounds (red lines) prescribed by the **d** static arena partitioning, **e** leader-follower and **f** peer-to-peer herding strategies (Color figure online)

performed. The movies of two illustrative trials are available in the Supplementary Material.⁵

Note that the dynamic control strategies we proposed were also found to have good performance when implemented in real robots with constraints imposed on their maximum velocities.

7 Conclusions

We presented a control strategy to solve the herding problem in the scenario where a group of herders is chasing a group of stochastic passive agents. Our approach is based on the combination of a set of local rules driving the herders according to the targets' positions and a herding strategy through which

the plane is partitioned among the herders, who then select the target to chase in the sector assigned to them either statically or dynamically. Our results show the effectiveness of the proposed strategy and its ability to cope with an increasing number of passive agents and variations of the repulsive force they feel when the herders approach them. Finally, we tested our control strategy via simulations in ROS and experiments on real robots, showing that our herding solutions are effective and viable in more realistic scenarios.

We wish to emphasise that to date our approach is the only one available in the literature to drive multiple herders to collect and contain a group of multiple agents that do not possess a tendency to flock and whose dynamics is stochastic. A pressing open problem is to derive a formal proof of convergence of the overall control system.

⁵ Movies available at <https://github.com/diBernardoGroup/HerdingProblem>.

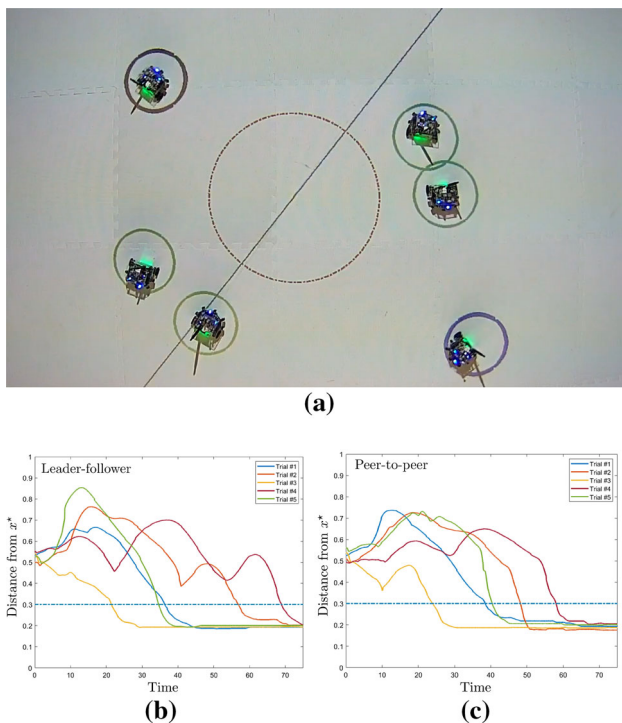


Fig. 7 Robotarium experiments. **a** Overview of the Robotarium arena, with GRITSBot robots for a herding scenario with $N_H = 2$ herder robots and $N_T = 4$ passive robots, and the evolution in time of the distance of the farthest passive robot from the containment region \mathcal{G} when the **b** leader-follower and **c** peer-to-peer strategies are employed for each trial. The mobile GRITSBots representing passive agents were initialised with random initial position chosen as $\mathbf{x}_i(0) = 2r^*e^{j\phi_i(0)}$, with $\phi_i(0)$ drawn with uniform distribution from the interval $(-\pi, \pi]$. The radius r^* of the containment region has been chosen equal to 0.3 m, i.e., equal to a third of the length of the arena

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10514-021-10033-6>.

Acknowledgements The authors wish to thank Dr. Jonathan Cacace from the University of Naples, Italy for his support with ROS.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A: Performance metrics

Denote with $\mathcal{X}(t) := \{i : \|\mathbf{x}_i(t) - \mathbf{x}^*\| \leq r^*\}$ the set of passive agents which are contained within the goal region \mathcal{G} at time t . Moreover, denote with $[0, T]$ the time interval over which the performance metrics are evaluated. The following metrics are used in the paper to evaluate the proposed herding strategies.

Gathering time defined as the time instant $t_g \in [0, T]$ such that condition (2) holds, that is, all the passive agents are in the containment region for all $t \geq t_g$.

Distance travelled by the herders which measures the mean in time and among herders of the distance travelled by the herders during the time interval $[0, t]$. It is defined as

$$d(t) := \frac{1}{N_H} \sum_{j=1}^{N_H} \frac{1}{t} \left(\int_0^t \|\dot{\mathbf{y}}_j(\tau)\| d\tau \right).$$

Therefore, $d_g := d(t_g)$, and $d_{\text{tot}} := d(T)$. A smaller average distance travelled indicates better efficiency of the herders in solving the task.

Herd distance from containment region which measures the herders ability to keep the herd close to the containment region, with centre \mathbf{x}^* . It is defined as the mean in time of the Euclidean distance between the centre of mass of the herd and the centre of the containment region, that is

$$D_T := \frac{1}{T} \int_0^T \left\| \left(\frac{1}{N_T} \sum_{i=1}^{N_T} \mathbf{x}_i(\tau) \right) - \mathbf{x}^*(\tau) \right\| d\tau.$$

A smaller average distance indicates better ability of the herders to keep the herd close to the containment region.

Herd spread measuring how much scattered the herd is in the game field. Denote as $\text{Pol}(t)$ the convex polygon defined by the convex hull of the points \mathbf{x}_i at time t , that is, $\text{Pol}(t) := \text{Conv}(\{\mathbf{x}_i(t), i = 1, \dots, N_T\})$. Then, the herd spread S is defined as the mean in time of the area of this polygon, that is

$$S := \frac{1}{T} \int_0^T \left(\int_{\text{Pol}(\tau)} d\mathbf{x} \right) d\tau.$$

Lower values corresponds to a more cohesive herd and consequently better herding performance. The herd spread can also be evaluated with respect to the area of the containment region, $A_{\text{cr}} = \pi(r^*)^2$, as $S_{\%} = S/A_{\text{cr}} \cdot 100$.

Appendix B: MATLAB simulations

In all simulations we considered the case of $N_H = 2$ or $N_H = 3$ artificial herders and $N_T = 7$ passive agents. Moreover, we considered a circular containment region with radius $r^* = 1$, centred in $\mathbf{x}^* = \mathbf{0}$, and a buffer region of width $\Delta r^* = 1$.

The numerical integration of the differential equations describing the dynamics of passive agents and herders has been realised using Euler–Maruyama method (Higham 2001) in the time interval $[0, T] = [0, 100]$ s with step size $dt = 0.006$ s, while the herder agents compute their next target-to-be-chased each $t_{\text{dwell}} = 50 dt$ s.

The values of all parameters used in the simulations were chosen as in Nalepka et al. (2017b). Collision detection radius $r_c = 0.0001$, coefficients of diffusion and repulsive motion $(\alpha_b, \alpha_r) = (0.005, 1)$, radial damping and stiffness coefficients $(b_r, \epsilon_r) = (11, 98.7)$, angular damping and stiffness coefficients $(b_\theta, \epsilon_\theta) = (11, 62.6)$.

The initial positions of the passive agents have been set outside the containment region as $\mathbf{x}_i(0) = 2r^*e^{j\phi_i(0)}$, $\forall i = 1, \dots, N_T$, with $\phi_i(0)$ drawn with uniform distribution in the interval $(-\pi, \pi]$, while the initial positions of herders have been taken on the circle with radius $4r^*$ and with angular displacement $(2\pi)/N_H$, to exclude from the simulations the rearrangement of the herders otherwise occurring at the beginning of each trial. Furthermore, collision avoidance forces between passive agents was also considered in the numerical simulations.

Specifically, the model (3) is extended by adding the term $\mathbf{V}_{c,i}(t)dt$, with

$$\mathbf{V}_{c,i}(t) = \sum_{i' \in \mathcal{X}_{c,i}(t)} \frac{\mathbf{x}_{i'}(t) - \mathbf{x}_i(t)}{\|\mathbf{x}_{i'}(t) - \mathbf{x}_i(t)\|^3},$$

where $\mathcal{X}_{c,i} := \{i' : \|\mathbf{x}_{i'} - \mathbf{x}_i\| \leq r_c\}$ is the set of all passive agents at time t inside the closed ball centred in \mathbf{x}_i with radius r_c .

Appendix C: ROS simulations

The mobile robots used for both passive and herder agents have been designed as Pioneer 3-DX robots driven by the differential drive controller provided in the set of ROS packages (`gazebo-ros-pkgs`) that allows the integration of Gazebo and ROS.

The environment and the robots share information through an exchange of messages that occurs publishing and subscribing to one or more of the available topics. A ROS node is attached to each herder and passive robots. It subscribes to the `/odom` topic; implements the agent's dynamics; and publishes a personalised `/cmd_vel` topic. The passive agents

collect odometric information from all the herders in the environment. The herder agents subscribe to the ID of the passive agent to-be-chased and collect its position. The published message is a velocity control input w.r.t. the robot's reference system to the differential drive of the robot: a translation v along x -axis and a rotation ω around z -axis of the robot. The reference trajectory $\mathbf{y}^*(t) = [r^* \cos \theta^*, r^* \sin \theta^*]^\top$, generated as in Sects. 3–4, is followed by each robot by means of the Cartesian regulator

$$\begin{aligned} v &= -k_v(\mathbf{y} - \mathbf{y}^*) [\cos \Phi, \sin \Phi] \\ \omega &= k_\omega(\theta^* - \Phi + \pi) \end{aligned}$$

where $\Phi(t)$ denotes the robot orientation w.r.t. the global reference system. The gains $k_v = 0.125$ and $k_\omega = 0.25$ have been tuned by trial-and-error to achieve smooth robot movements.

The initial position of the agents have been set as in Appendix B.

The target selection strategies (Sect. 4.2) are processed in an *ad hoc* ROS node. It subscribes to the odometry topic; computes the user-chosen strategy (i.e. global, static arena partitioning, leader-follower or peer-to-peer); and publishes a custom message with the ID of the targets to-be-chased on the `/herder/chased_target` topic. The custom message is an array of integer numbers, its j -th element corresponds to the passive agent chased by the j -th herder robot.

The Gazebo-ROS simulations were run on Ubuntu 18.0404 LTS hosted on a VirtualMachine with a 10GB RAM with ROS Melodic distribution and Gazebo 9.13.0.

References

- Auletta, F., di Bernardo, M., & Richardson, M. J. (2021). Human-inspired strategies to solve complex joint tasks in multi agent systems, *IFAC-PapersOnLine*, 54, pp. 105–110.
- Bürger, M., Notarstefano, G., Allgöwer, F., & Bullo, F. (2011). A distributed simplex algorithm and the multi-agent assignment problem. In: *Proceedings of the American control conference* (pp. 2639–2644). <https://doi.org/10.1109/acc.2011.5990932>.
- Chipade, V. S., & Panagou, D. (2019). Herding an adversarial swarm in an obstacle environment. In: *Proceedings of the IEEE conference on decision and control* (pp. 3685–3690). IEEE, <https://doi.org/10.1109/CDC40024.2019.9029573>.
- Chipade, V. S., Marella, V. S. A., & Panagou, D. (2021). Aerial swarm defense by StringNet herding: Theory and experiments. *Frontiers in Robotics and AI*, 8, 81. <https://doi.org/10.3389/frobt.2021.640446>
- Haque, M., Rahmani, A., & Egerstedt, M. (2009). A hybrid, multi-agent model of foraging bottlenose dolphins. *IFAC Proceedings Volumes*, 42(17), 262–267. <https://doi.org/10.3182/20090916-3-ES-3003.00046> 3rd IFAC conference on analysis and design of hybrid systems.
- Haque, M. A., Rahmani, A. R., & Egerstedt, M. B. (2011). Biologically inspired confinement of multi-robot systems. *International Jour-*

- nal of Bio-Inspired Computation*, 3(4), 213–224. <https://doi.org/10.1504/IJBIC.2011.041145>
- Higham, D. J. (2001). An algorithmic introduction to numerical simulation of stochastic differential equations. *SIAM Review*, 43(3), 525–546. <https://doi.org/10.1137/S0036144500378302>
- Lee, W., & Kim, D. E. (2017). Autonomous herding behaviors of multiple target steering robots. *Sensors*, 17(12), 2729. <https://doi.org/10.3390/s17122729>
- Licitra, R. A., Hutcheson, Z. D., Doucette, E. A., & Dixon, W. E. (2017). Single agent herding of n-agents: A switched systems approach. *IFAC-PapersOnLine*, 50(1), 14374–14379. <https://doi.org/10.1016/j.ifacol.2017.08.2020> 20th IFAC world congress.
- Licitra, R. A., Bell, Z. I., Doucette, E. A., & Dixon, W. E. (2018). Single agent indirect herding of multiple targets: A switched adaptive control approach. *IEEE Control Systems Letters*, 2(1), 127–132. <https://doi.org/10.1109/LCSYS.2017.2763968>
- Licitra, R. A., Bell, Z. I., & Dixon, W. E. (2019). Single-agent indirect herding of multiple targets with uncertain dynamics. *IEEE Transactions on Robotics*, 35(4), 847–860. <https://doi.org/10.1109/TRO.2019.2911799>
- Lien, J. M., Bayazit, O., Sowell, R., Rodriguez, S., & Amato, N. M. (2004). Shepherding behaviors. In: *Proceedings of the IEEE international conference on robotics and automation* (pp. 4159–4164). <https://doi.org/10.1109/ROBOT.2004.1308924>.
- Lien, J. M., Rodriguez, S., Malric, J., & Amato, N. (2005). Shepherding behaviors with multiple shepherds. In: *Proceedings of the IEEE international conference on robotics and automation* (pp. 3402–3407). <https://doi.org/10.1109/ROBOT.2005.1570636>.
- Long, N. K., Sammut, K., Sgarioni, D., Garratt, M., & Abbass, H. A. (2020). A comprehensive review of herding as a bio-inspired swarm-robotics guidance approach. *IEEE Transaction on Emerging Topics in Computational Intelligence*, 4(4), 523–537. <https://doi.org/10.1109/TETCI.2020.2992778>
- Montijano, E., Priolo, A., Gasparri, A., & Sagues, C. (2013). Distributed entrapment for multi-robot systems with uncertainties. In: *Proceedings of the IEEE conference on decision and control* (pp. 5403–5408). IEEE, <https://doi.org/10.1109/CDC.2013.6760739>.
- Nalepka, P., Riehm, C., Mansour, C. B., Chemero, A., & Richardson, M. J. (2015). Investigating strategy discovery and coordination in a novel virtual sheep herding game among dyads. In: *Proceedings of the 37th Annual Meeting of the Cognitive Science Society* (pp. 1703–1708).
- Nalepka, P., Kallen, R. W., Chemero, A., Saltzman, E., & Richardson, M. J. (2017). Herd those sheep: Emergent multiagent coordination and behavioral-mode switching. *Psychological Science*, 28(5), 630–650. <https://doi.org/10.1177/0956797617692107>
- Nalepka, P., Lamb, M., Kallen, R. W., Saltzman, E., Chemero, A., & Richardson, M. J. (2017b). First step is to group them: Task-dynamic model validation for human multiagent herding in a less constrained task. In: *Proceedings of the 39th Annual Meeting of the Cognitive Science Society* (pp. 2784–2789).
- Nalepka, P., Lamb, M., Kallen, R. W., Shockley, K., Chemero, A., Saltzman, E., & Richardson, M. J. (2019). Human social motor solutions for human-machine interaction in dynamical task contexts. *PNAS*, 116(4), 1437–1446. <https://doi.org/10.1073/pnas.1813164116>
- Nolfi, S. (2002). Power and the limits of reactive agents. *Neurocomputing*, 42(1–4), 119–145. [https://doi.org/10.1016/S0925-2312\(01\)00598-7](https://doi.org/10.1016/S0925-2312(01)00598-7)
- Paranjape, A. A., Chung, S., Kim, K., & Shim, D. H. (2018). Robotic herding of a flock of birds using an unmanned aerial vehicle. *IEEE Transactions on Robotics*, 34(4), 901–915. <https://doi.org/10.1109/TRO.2018.2853610>
- Pickem, D., Lee, M., Egerstedt, M. (2015). The GRITSBot in its natural habitat—a multi-robot testbed. In: *Proceedings of the IEEE international conference on robotics and automation* (pp. 4062–4067). <https://doi.org/10.1109/ICRA.2015.7139767>.
- Pickem, D., Glotfelter, P., Wang, L., Mote, M., Ames, A., Feron, E., & Egerstedt, M. (2017). The Robotarium: A remotely accessible swarm robotics research testbed. In: *Proceedings of the IEEE international conference on robotics and automation* (pp. 1699–1706). <https://doi.org/10.1109/ICRA.2017.7989200>.
- Pierson, A., & Schwager, M. (2018). Controlling noncooperative herds with robotic herders. *IEEE Transaction on Robotics*, 34(2), 517–525. <https://doi.org/10.1109/TRO.2017.2776308>
- Rigoli, L. M., Nalepka, P., Douglas, H., Kallen, R. W., Hosking, S., Best, C., Saltzman, E., & Richardson, M. J. (2020). Employing models of human social motor behavior for artificial agent trainers. In: *Proceedings of the 19th international conference on autonomous agents and multiAgent systems* (pp. 1134–1142).
- Sebastián, E., & Montijano, E. (2021). Multi-robot implicit control of herds. In: *Proceedings of the IEEE international conference on robotics and automation* (To appear).
- Song, H., Varava, A., Kravchenko, O., Kragic, D., Wang, M. Y., Pokorný, F. T., & Hang, K. (2021). Herding by caging: A formation-based motion planning framework for guiding mobile agents. *Autonomous Robots*, 45, 613–631. <https://doi.org/10.1007/s10514-021-09975-8>
- Strombom, D., Mann, R. P., Wilson, A. M., Hailes, S., Morton, A. J., Sumpter, D. J., & King, A. J. (2014). Solving the herding problem: Heuristics for herding autonomous, interacting agents. *Journal of the Royal Society Interface*, 11, 20140719. <https://doi.org/10.1098/rsif.2014.0719>
- Varava, A., Hang, K., Kragic, D., & Pokorný, F. T. (2017). Herding by caging: A topological approach towards guiding moving agents via mobile robots. In: *Proceedings of the robotics: Science and systems XIII*. <https://doi.org/10.15607/rss.2017.xiii.074>.
- Vaughan, R., Sumpter, N., Henderson, J., Frost, A., & Cameron, S. (2000). Experiments in automatic flock control. *Robotics and Autonomous Systems*, 31(1), 109–117. [https://doi.org/10.1016/S0921-8890\(99\)00084-6](https://doi.org/10.1016/S0921-8890(99)00084-6)
- Wilmarth, S. A., Amato, N. M., & Stiller, P. F. (1999). MAPRM: A probabilistic roadmap planner with sampling on the medial axis of the free space. In: *Proceedings of the IEEE international conference on robotics and automation* vol. 2 (pp. 1024–1031). <https://doi.org/10.1109/ROBOT.1999.772448>
- Wilson, S., Glotfelter, P., Wang, L., Mayya, S., Notomista, G., Mote, M., & Egerstedt, M. (2020). The Robotarium: Globally impactful opportunities, challenges, and lessons learned in remote-access, distributed control of multirobot systems. *IEEE Control Systems Magazine*, 40(1), 26–44. <https://doi.org/10.1109/MCS.2019.2949973>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Fabrizia Auletta attended University of Naples Federico II where she received her Bachelor degree in Electronic Engineering. In 2018 she obtained her Master degree in Engineering of Automation. Her Master project was carried out in collaboration with Professor John Hogan at the University of Bristol. She is now enrolled in a Cotutelle PhD programme between the University of Bristol (UK) and Macquarie University (Australia), working with Professors Mario di Bernardo and Michael J. Richard-

son. Her research focuses on the application of control theory and machine learning techniques to study and design autonomous agents in complex joint tasks.



Davide Fiore is Assistant Professor of Automatic Control at the Department of Mathematics and Applications "R. Caccioppoli", University of Naples Federico II. He is currently working on synthetic biological applications. He received his Bachelor's and Master's degrees in Automation Engineering from University of Naples Federico II in 2010 and 2013, respectively. In May 2017, he received his Ph.D. in Information Technology and Electrical Engineering from the University of

Naples Federico II, Italy, under the supervision of Prof. Mario di Bernardo. From 2017 to 2019 he held a Postdoctoral Researcher position at the University of Naples Federico II working on convergence analysis and control of stochastic multi-agent systems in synthetic biological applications. From May 2020 he is Assistant Professor—fixed term, at the University of Naples Federico II, Italy. His research interests include nonlinear systems and control, switched and hybrid systems, multi-agent systems, and synthetic biology.



Michael J. Richardson is Professor in the Department of Psychology at Macquarie University in Sydney, Australia. His research is directed towards understanding the lawful dynamics of human perception, action, and cognition. He has expertise in experimental and applied psychology, cognitive science, human-movement science,

perception-action, joint-action and social coordination, virtual-reality, complex systems, quantitative and statistical analysis methods, and

dynamical modelling. He teaches a range of graduate and undergraduate courses on these topics, as well as workshops on nonlinear time-series analysis and dynamical systems modelling.



Mario di Bernardo (SMIEEE '06, FIEEE 2012) is Professor of Automatic Control at the University of Naples Federico II, Italy and Visiting Professor of Nonlinear Systems and Control at the University of Bristol, U.K. He currently serves as Deputy pro-Vice Chancellor for Internationalization at the University of Naples and coordinates the research area on Modeling and Engineering Risk and Complexity of the Scuola Superiore Meridionale, the new School of Advanced Studies

located in Naples. On 28th February 2007 he was bestowed the title

of Cavaliere of the Order of Merit of the Italian Republic for scientific merits from the President of Italy. He was elevated to the grade of Fellow of the IEEE in January 2012 for his contributions to the analysis, control and applications of nonlinear systems and complex networks. In 2009, he was elected President of the Italian Society for Chaos and Complexity for the term 2010–2013. He was re-elected in 2013 for the term 2014–2017. In 2006 and again in 2009 he was elected to the Board of Governors of the IEEE Circuits and Systems Society. From 2011 to 2014 he was Vice President for Financial Activities of the IEEE Circuits and Systems Society. In 2015 he was appointed to the Board of Governors of the IEEE Control Systems Society. He was Distinguished Lecturer of the IEEE Circuits and Systems Society for the two-year term 2016–2017. His research interests include the analysis, synchronization and control of complex network systems; piecewise-smooth dynamical systems; nonlinear dynamics and nonlinear control with applications to engineering and computational biology. He authored or co-authored more than 220 international scientific publications including more than 110 papers in scientific journals, a research monograph and two edited books. According to the international database SCOPUS (January 2022), his h-index is 50 and his publications received over 10000 citations by other authors. In 2017, he received the IEEE George N. Saridis Best Transactions Paper Award for Outstanding Research. He serves on the Editorial Board of several international scientific journals and conferences. From 1st January 2014 till 31st December 2015 he was Deputy Editor-in-Chief of the IEEE Transactions on Circuits and Systems: Regular Papers. He is Senior Editor of the IEEE Transactions on Control of Network Systems and Associate Editor of the IEEE Control Systems Letters, the Conference Editorial Board of the IEEE Control System Society and the European Control Association (EUCA). He was Associate Editor of Nonlinear Analysis: Hybrid Systems; the IEEE Transactions on Circuits and Systems I: Regular Papers from 1999 to 2002 and again from 2008 to 2010, and the IEEE Transactions on Circuits and Systems II: Brief papers from 2003 till 2008. He is regularly invited as Plenary Speakers in Italy and abroad. He was Program co-Chair of the European Control Conference 2019, Publicity Chair of the IEEE ISCAS Conference 2018 and has been organizer and co-organizer of several scientific initiatives and events and received funding from several funding agencies and industry including the European Union, the UK research councils the Italian Ministry of Research and University.