Solving the multi-agent shepherding problem

harnessing complex systems for control

Mario di Bernardo

University of Naples Federico II & Scuola Superiore Meridionale, Italy

Joint IEEE CSS & CASS Chapters, Vancouver – 25th September 2024

Outline

- Controlling complex systems
- Shepherding as a paradigmatic control task
- A brief overview of existing solutions
- Removing some strong assumptions
- Herdability of a complex multiagent system
- A machine learning approach
- Conclusions, perspectives and applications

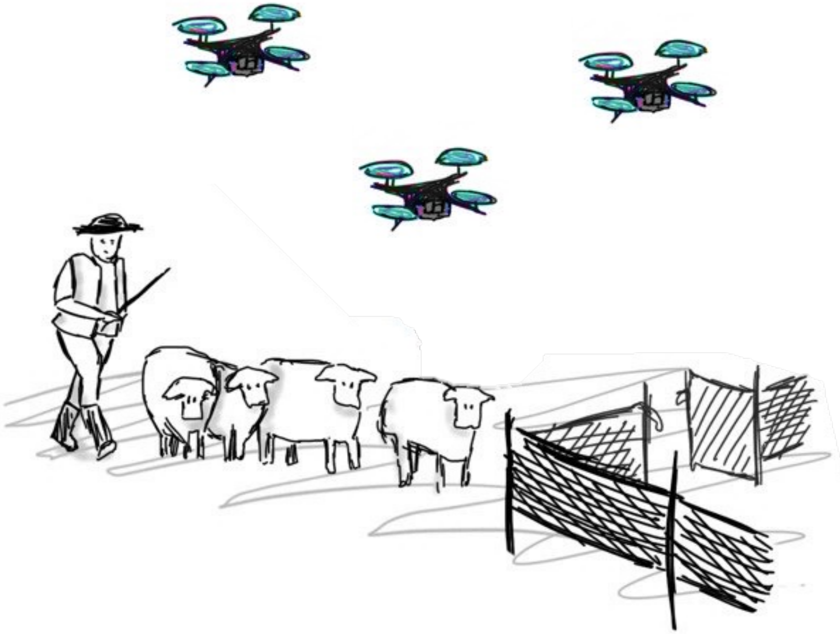

© F. Auletta

Taming complexity

- From power grids and swarm robotics to biology and epidemiology
- Often, we wish to control the emerging collective behaviour of complex system
- E.g. avoid or induce synchronization, pattern formation, prevent undesired cascading phenomena, achieve crowd control etc

Can we orchestrate in real-time the collective behaviour of a complex system?

Control and complex systems

- Attention has focussed on the problem of *controlling a complex network*..

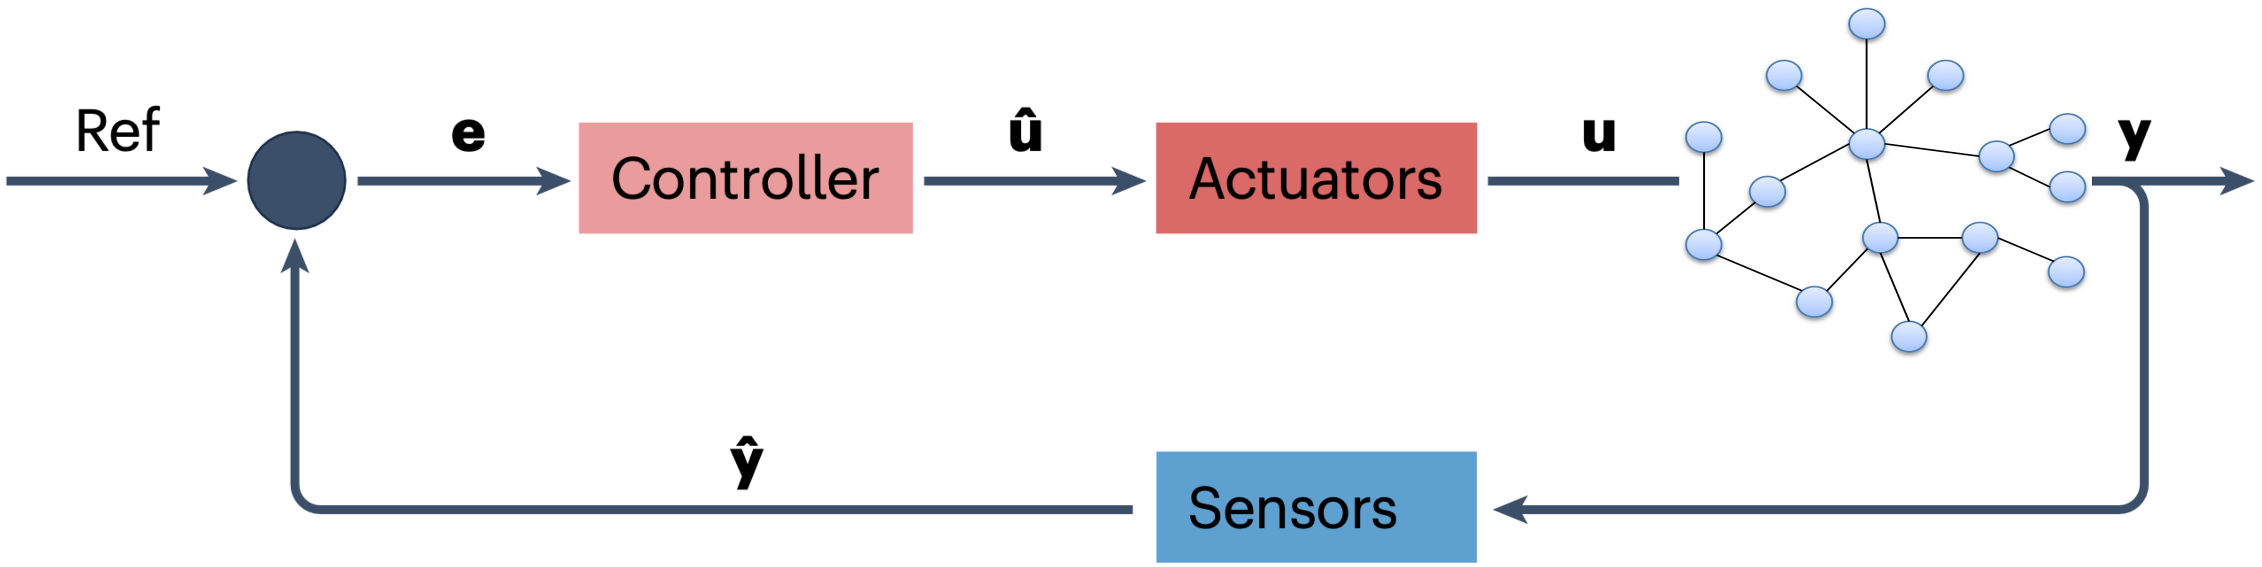

R. De Souza, MdB, YY Liu, "Controlling complex systems with complex nodes", Nature Reviews Physics, 2023

The key ingredients

- *Can we orchestrate in real-time the collective behaviour of a complex system?*

Feedback Control = Sense + Compute + Actuate

1. Whom do we sense? *observability*
2. Whom do we control? *controllability*
3. What do we compute? *controller design*

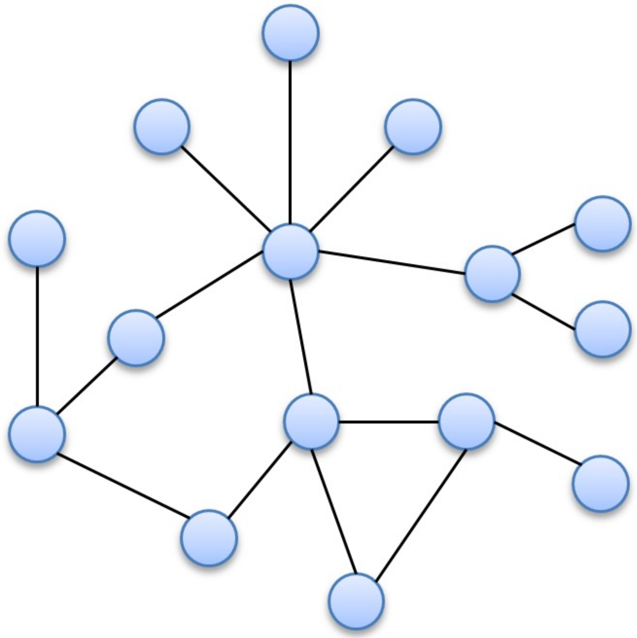

- We want the control strategy to be distributed and to be computed in real-time as a function of the sensed variables

Control and Complex Systems

- We need to ”close the loop” *across different scales*

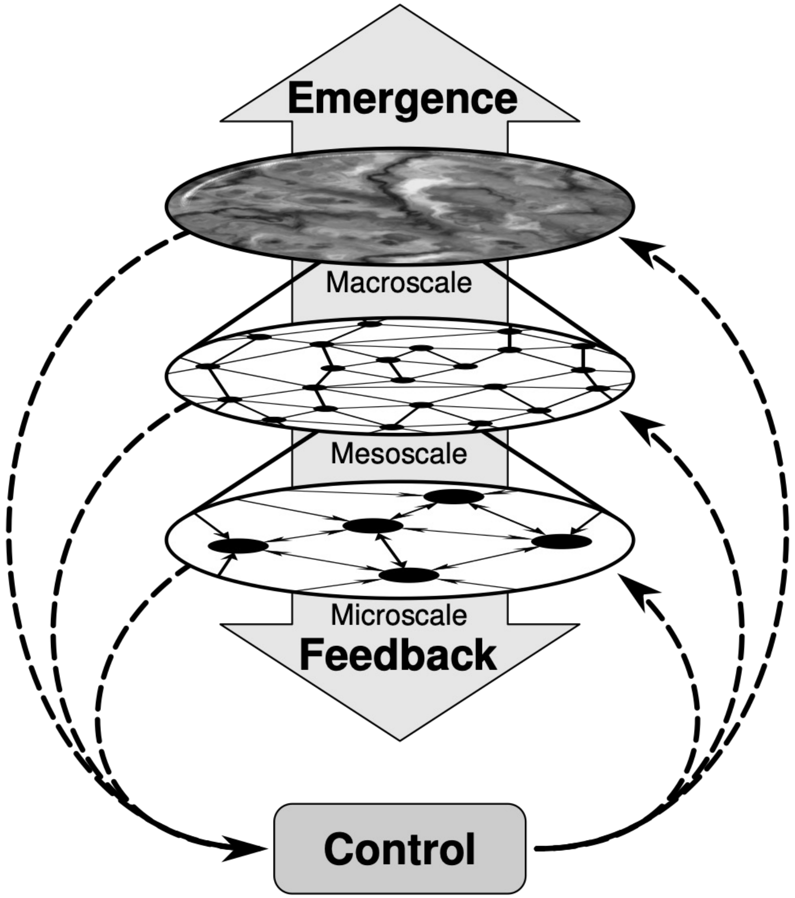

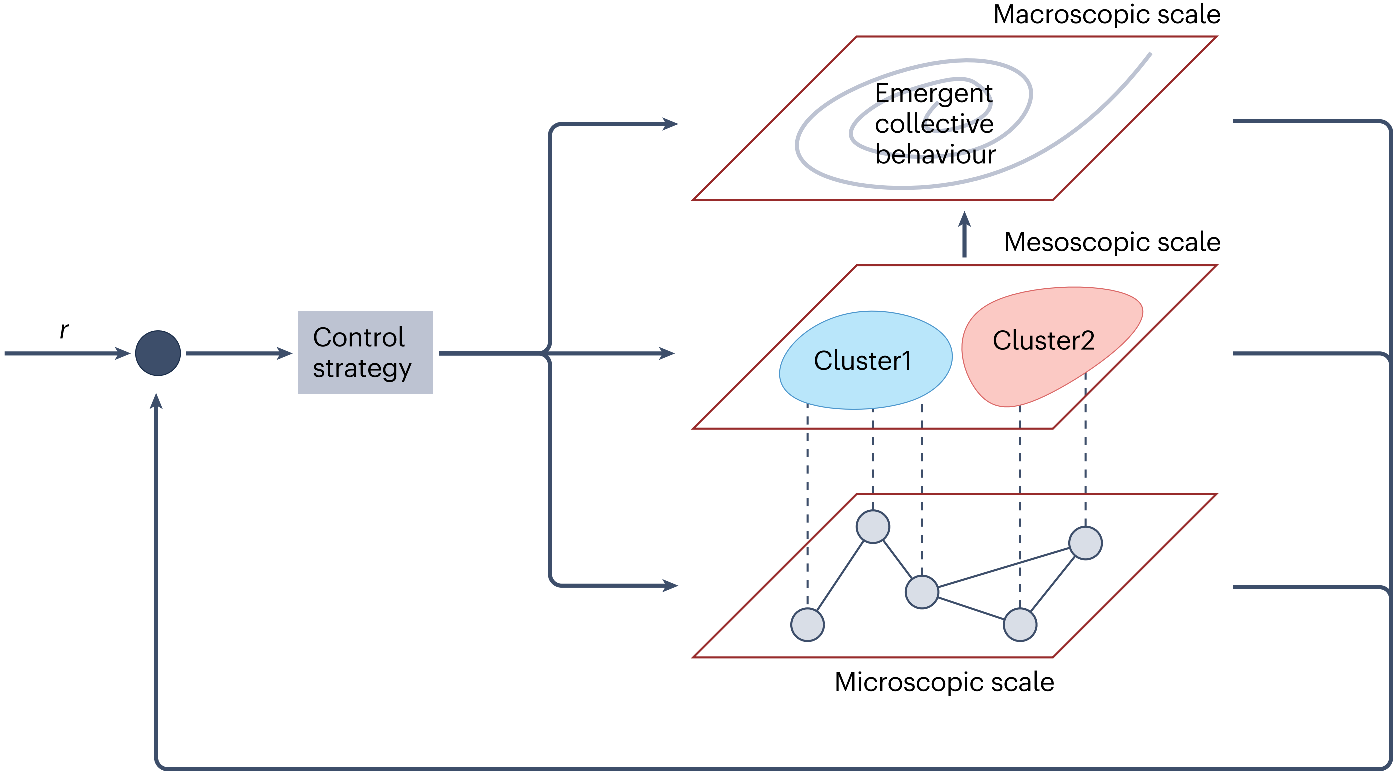

R. D’Souza, M. di Bernardo, YY Liu “Controlling complex networks with complex nodes”, Nature Reviews Physics, 2023

Complex systems for control

- What if the complex system acts as the controller rather than being the system we wish to control?
- *Can we "engineer" the collective behaviour of a complex system to perform a control task?*

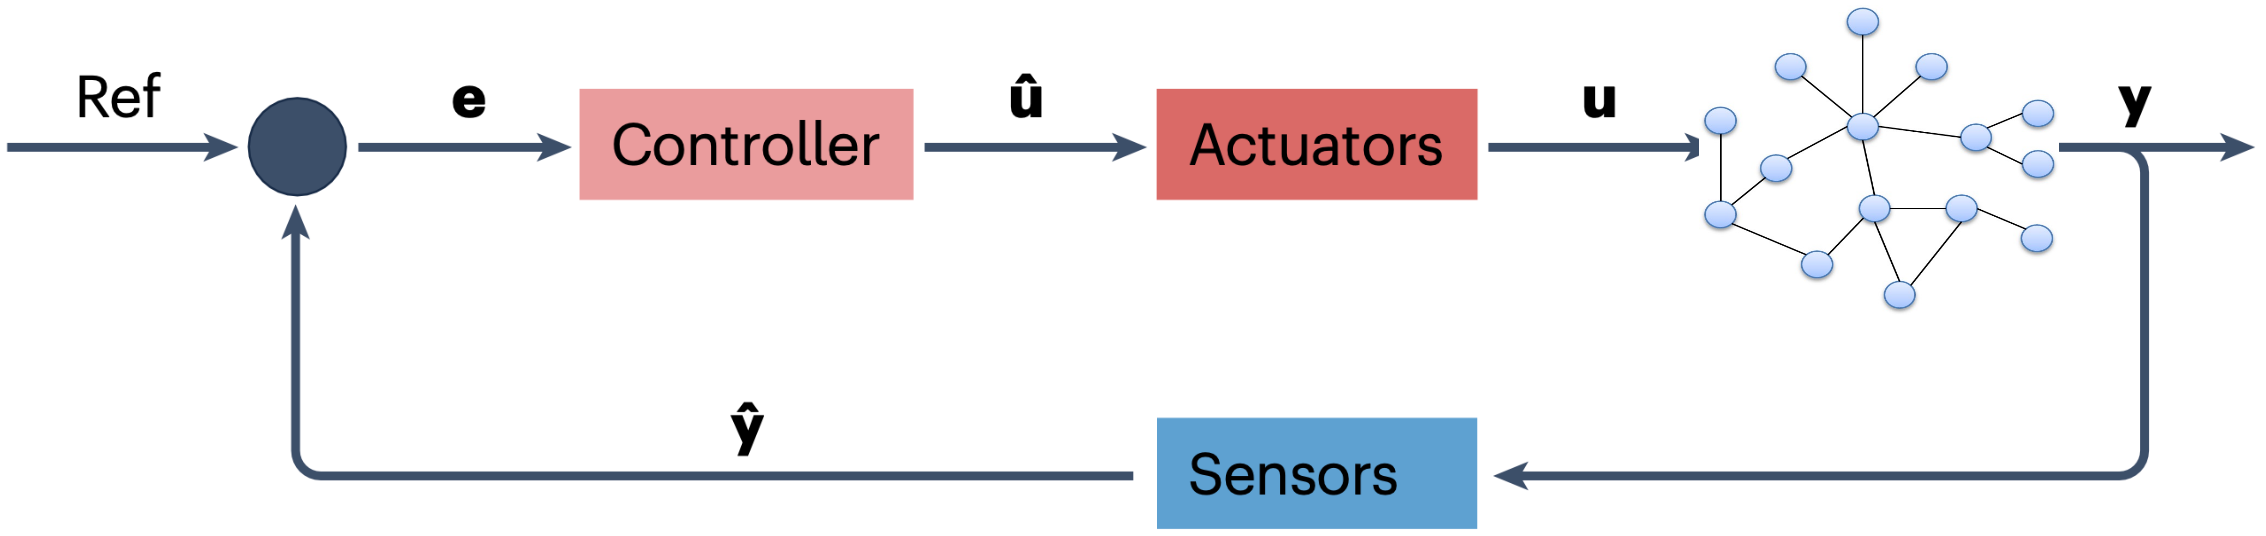

The shepherding control problem

- The shepherding problem is a paradigmatic example
- Here a group of agents, *the herders*, need to steer the collective dynamics of another group of agents, *the targets,* in some desired way

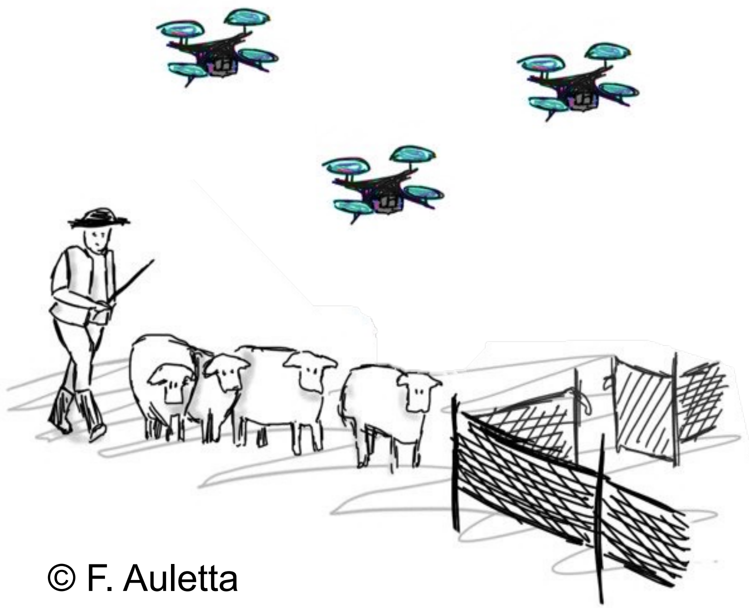

© F. Auletta

Relevance

- Observed in *biological systems* (e.g dolphins hunting fish [Haque et al, *2011,Int.* J. Bio-Inspired Comp], ants collecting aphids [Oliver et al,2007,Proc. R. Soc. B])
- *Technological applications*: search & rescue, crowd control, oil cleanup [Long et al, 2021, IEEE Emerging Comp applications]
- *Swarm robotics* and shepherding robots
- Active matter *physics* etc



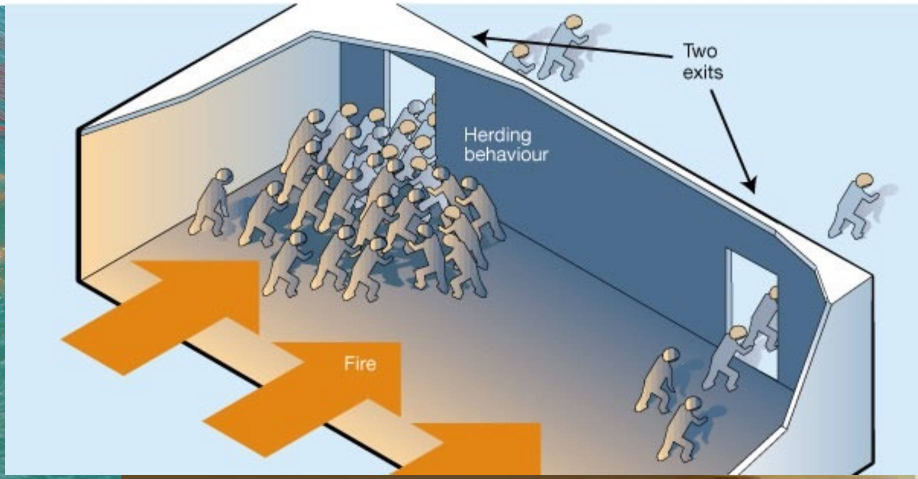

The planar shepherding problem

- A group of agents, *the herders*, is tasked with the goal of collecting and coralling another group of agents, *the targets* towards some desired goal region in the plane

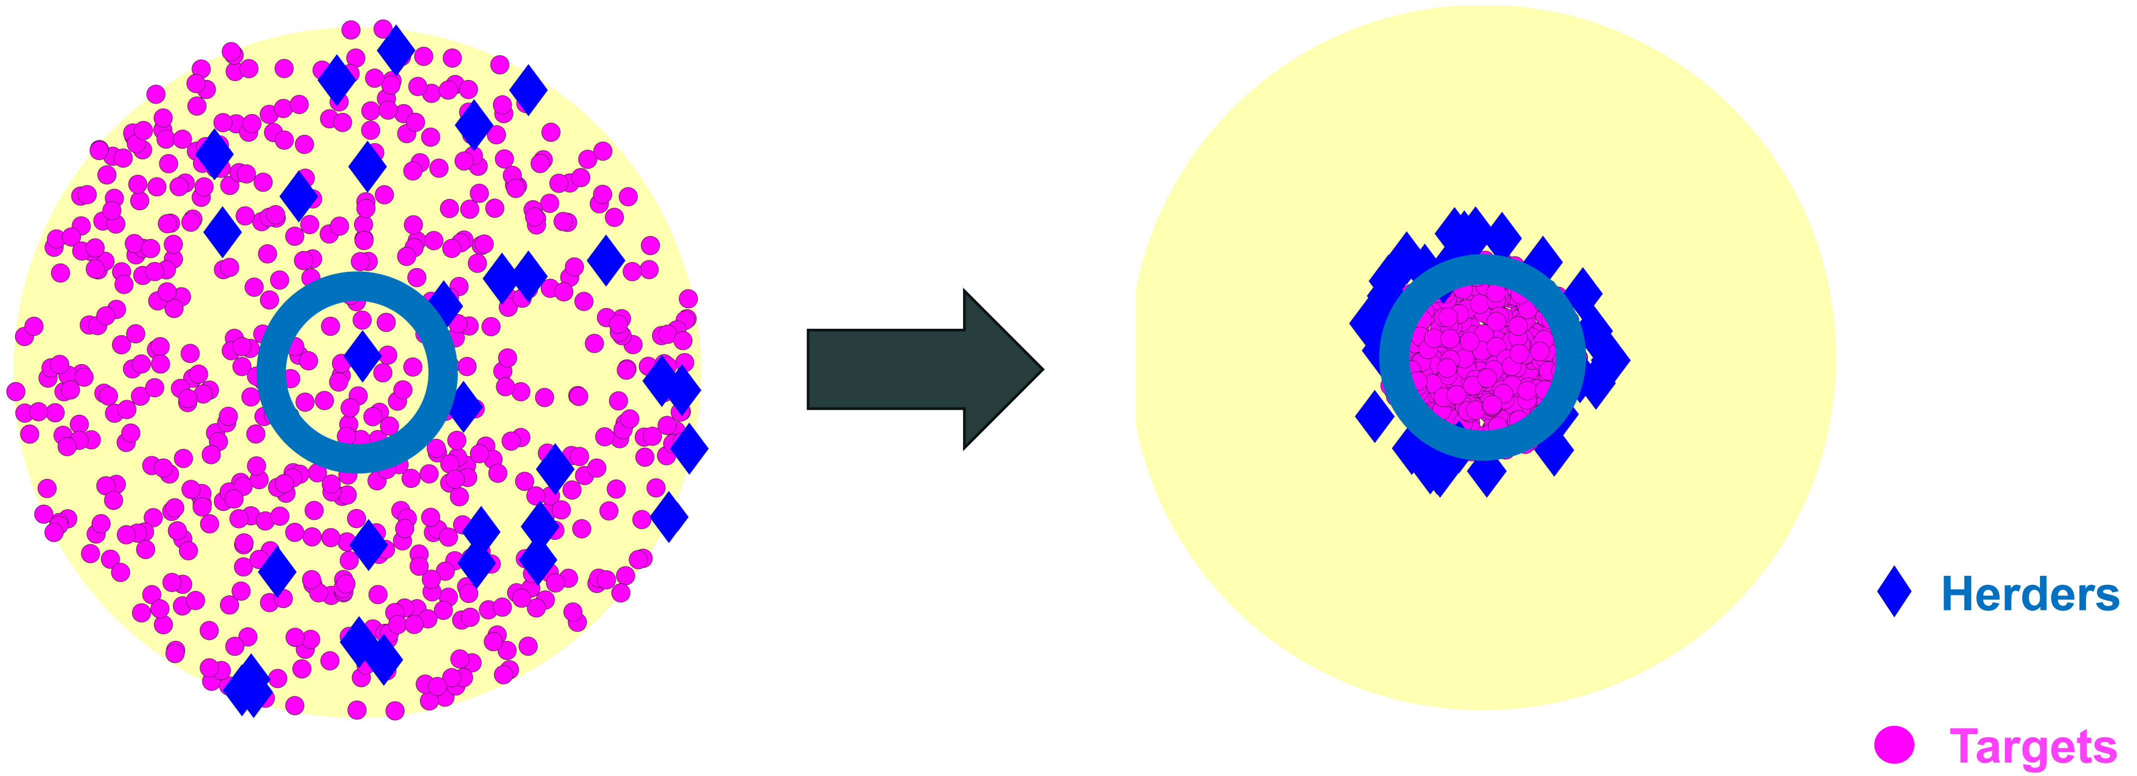

F. Auletta, D. Fiore et al , "Herding stochastic autonomous agents via local control rules and online global target selection strategies", *Autonomous Robots*, 2022

A complex system performing a control task

- In Shepherding the emerging collective behaviour of a complex system of *targets* must be controlled by driving the emerging behaviour of another complex system (the *herders*)
- A task also referred to as *indirect control* in the literature

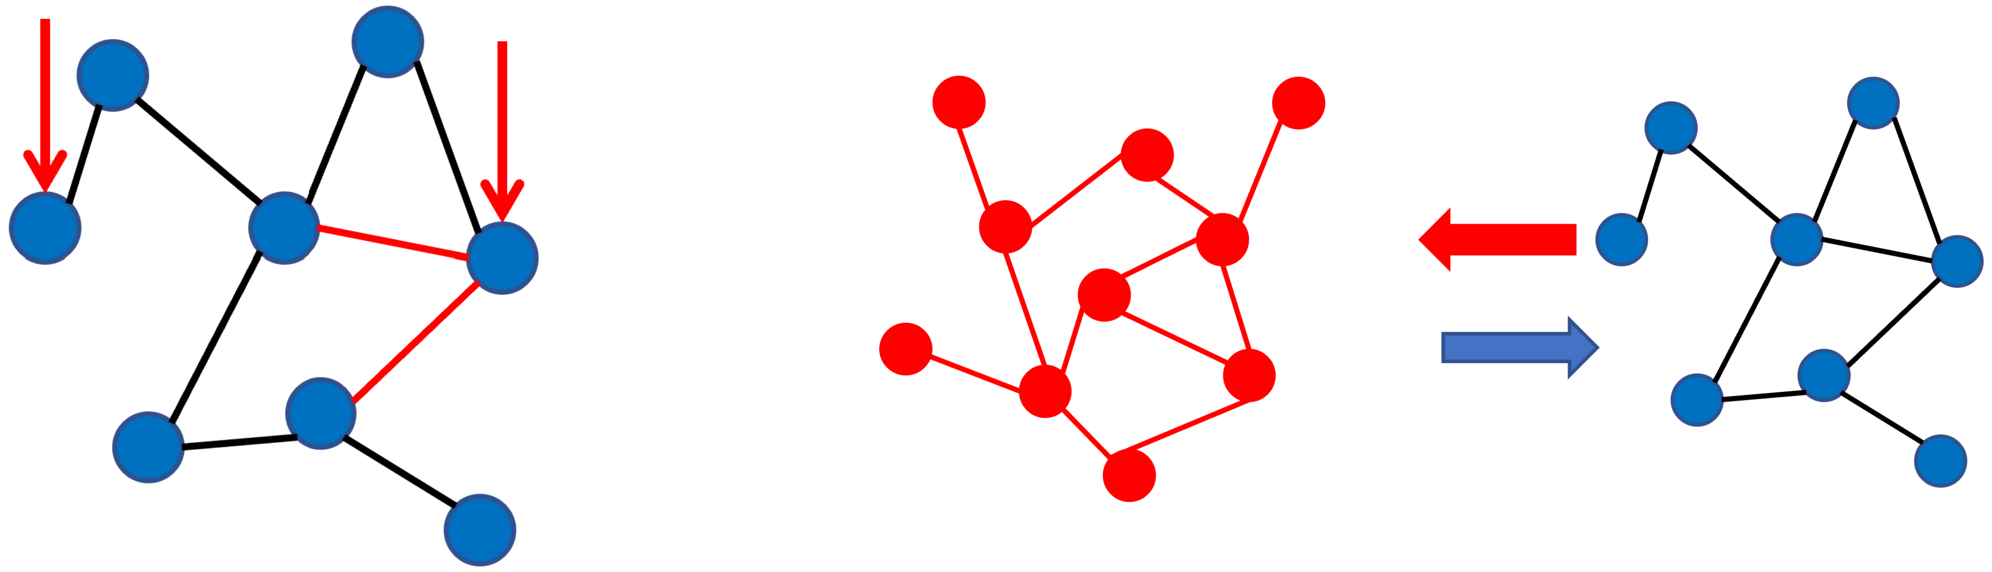

The shepherding control problem

- The crucial problem is the design of the herders' dynamics so as to achieve the desired goal
- *Herders* must steer the target behaviour towards the desired region (*coralling task)*
- Also, they need to *cooperate with each other* and collectively implement decision-making strategies (*target selection*)
- Herders can possess *global* or *local* information according to their sensing regions

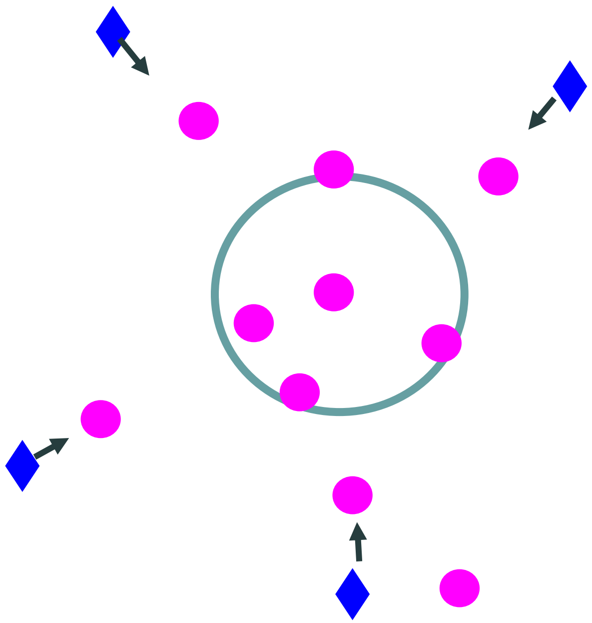

A Pierson, M Schwager, Controlling noncooperative herds with robotic herders, IEEE Trans Robotics 2018
R.A. Licitra, Z Bell, W Dixon, Single-agent indirect herding of mutliple targets with uncertain dynamics", IEEE Trans Robotics, 2019
D. Ko, E. Zuazua, Asymptotic behaviour and control of "a guidance by repulsion model", Math Models Methods Appl Sci, 2020

Targets behaviour

- Targets usually have their own dynamics (fish schooling, crowds, animal groups..)
- Their motion being influenced by the presence of an herder within their region of influence
- For instance, targets can be repelled from (or attracted to) nearby herders
- Typically, targets are assumed to **flock** together (selfish herd hypothesis)
- An assumption that is often unrealistic in some circumstances

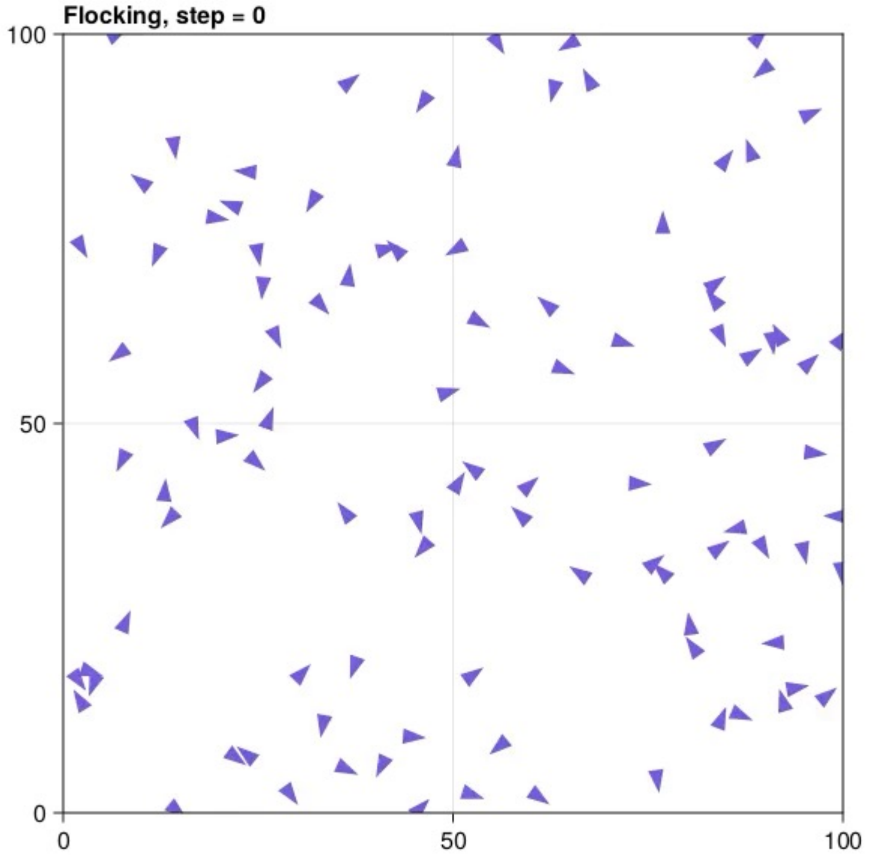

Existing solutions

- Many solutions are available in the case of 1 herder and 1 or many targets or when the number of herders is equal to the number of targets.
- For more general cases, some of the earliest solutions involve path planning via a global rule based approach (not easily scalable and computationally expensive!)

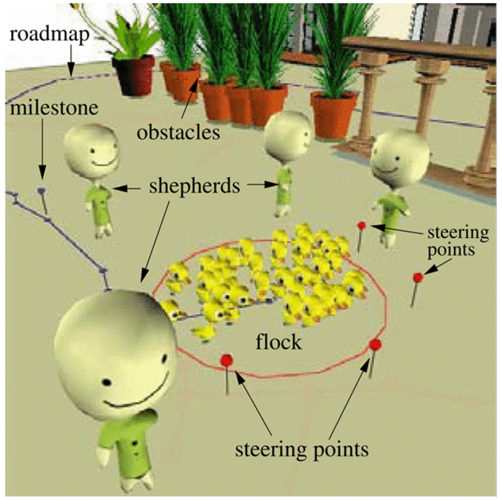

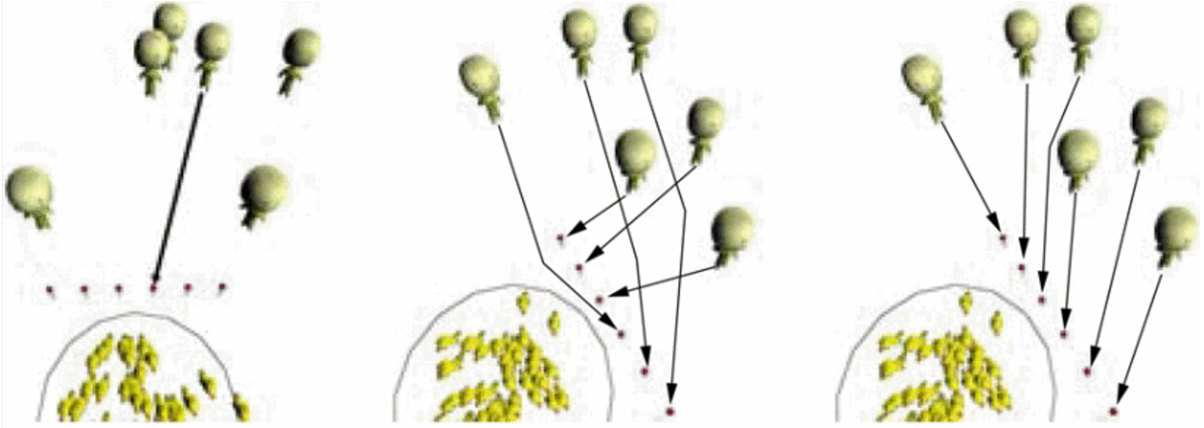

(a) (b) (c)

J.-M. Lien, S. Rodriguez, J. Malric, and N. Amato, *Shepherding Behaviors with Multiple Shepherds*, in IEEE International Conference on Robotics and Automation, 2005, pp. 3402–3407.

P. Kachroo, S. Shedied, J. Bay, and H. Vanlandingham, *Dynamic programming solution for a class of pursuit evasion problems: the herding problem*, IEEE Transactions on Systems, Man and Cybernetics, Part C (Applications and Reviews), 31 (2001), pp. 35–41.

Optimal “guidance by repulsion” models

- In the mathematics community the problem has been studied most notably by Zuazua and coworkers [2020]
- The herders (or drivers) know and track the barycenter of the **flocking** evaders
- No collaborative strategy is set up between the herders just avoiding collisions
- The herder dynamics is the *off-line* solution of open-loop optimal control problems

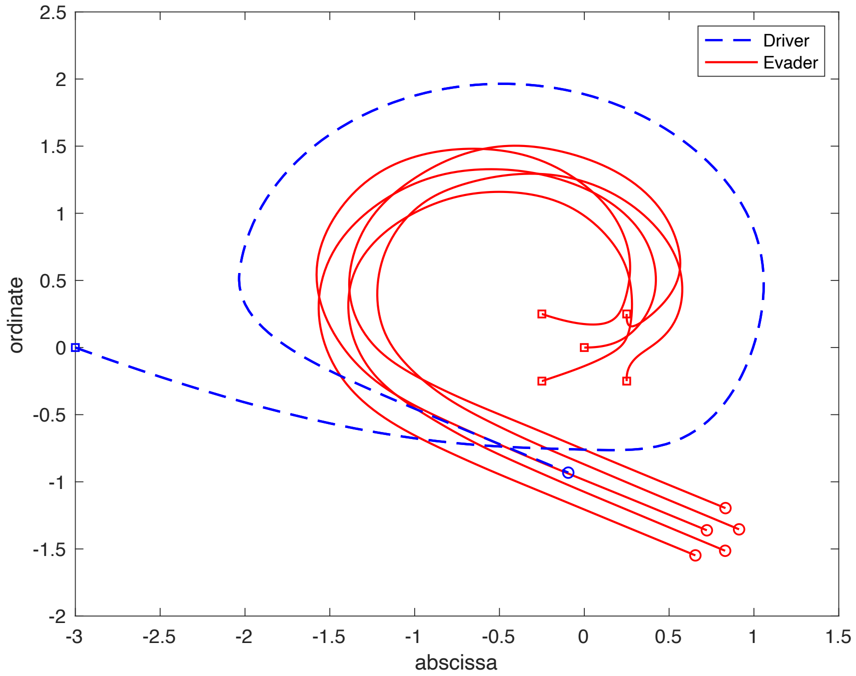

- Some feedback laws inspired by these solutions…
- .. but “*if the ensemble of evaders is separated and hard to flock together initially, then this strategy does not work*”

D. Ko, E. Zuazua, Asymptotic behaviour and control of “a guidance by repulsion model”, Math Models Methods Appl Sci, 2020

Bio-inspired models

- Other solutions aimed at replicating behaviour observed in natural systems (shepherds and sheepdogs, dolphins foraging etc) by assigning specific dynamics for targets and herders

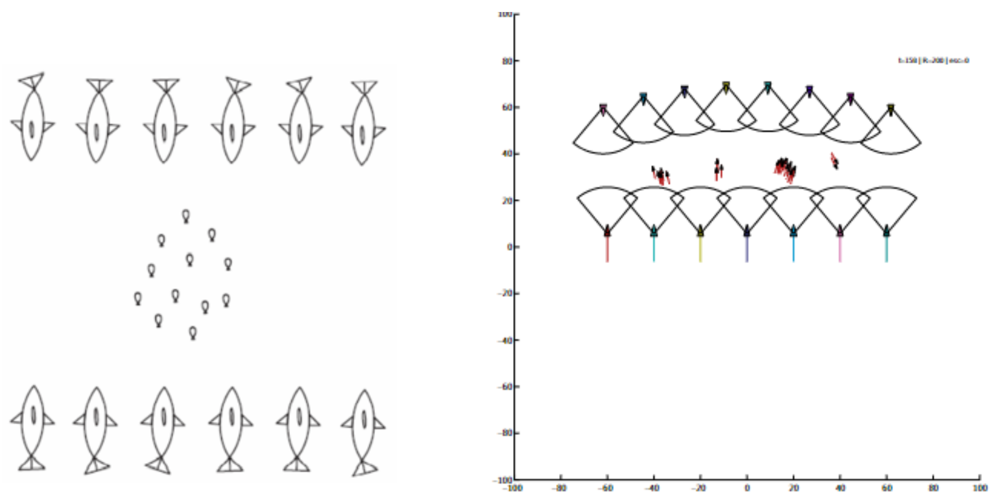

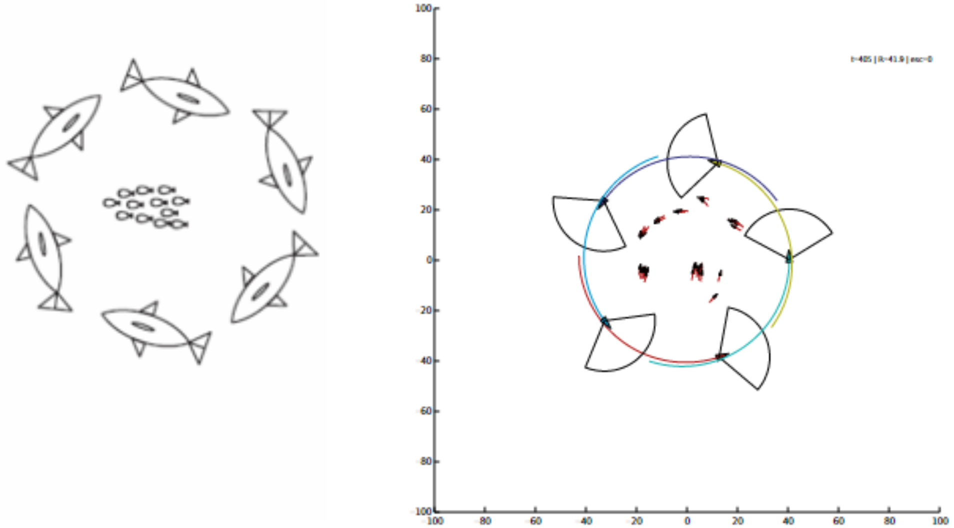

M. A. HAQUE, A. R. RAHMANI, AND M. B. EGERSTEDT, *Biologically inspired confinement of multi-robot systems*, International Journal of Bio-Inspired Computation, 3 (2011), pp. 213–224.

M. HAQUE, A. RAHMANI, AND M. EGERSTEDT, *A hybrid, multi-agent model of foraging bottlenose dolphins*, IFAC Proceedings Volumes, 42 (2009), pp. 262–267. 3rd IFAC Conference on Analysis and Design of Hybrid Systems.

Human-inspired models

- Recently, herding tasks have received growing attention in cognitive and psychological sciences to study human decision-making, complex joint action and team coordination

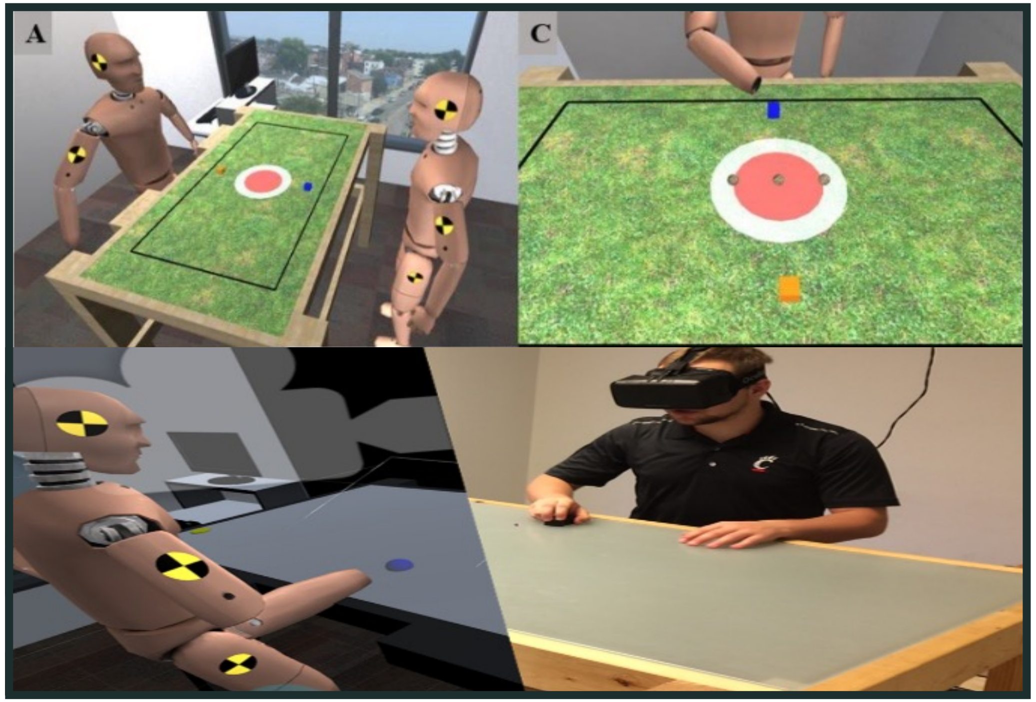

P. Nalepka, R. W. Kallen, A. Chemero, E. Saltzman, and M. J. Richardson, *Herd Those Sheep: Emergent Multiagent Coordination and Behavioral-Mode Switching*, Psychological Science, 28 (2017), pp. 630–650.

The “human” solution

Search and Rescue

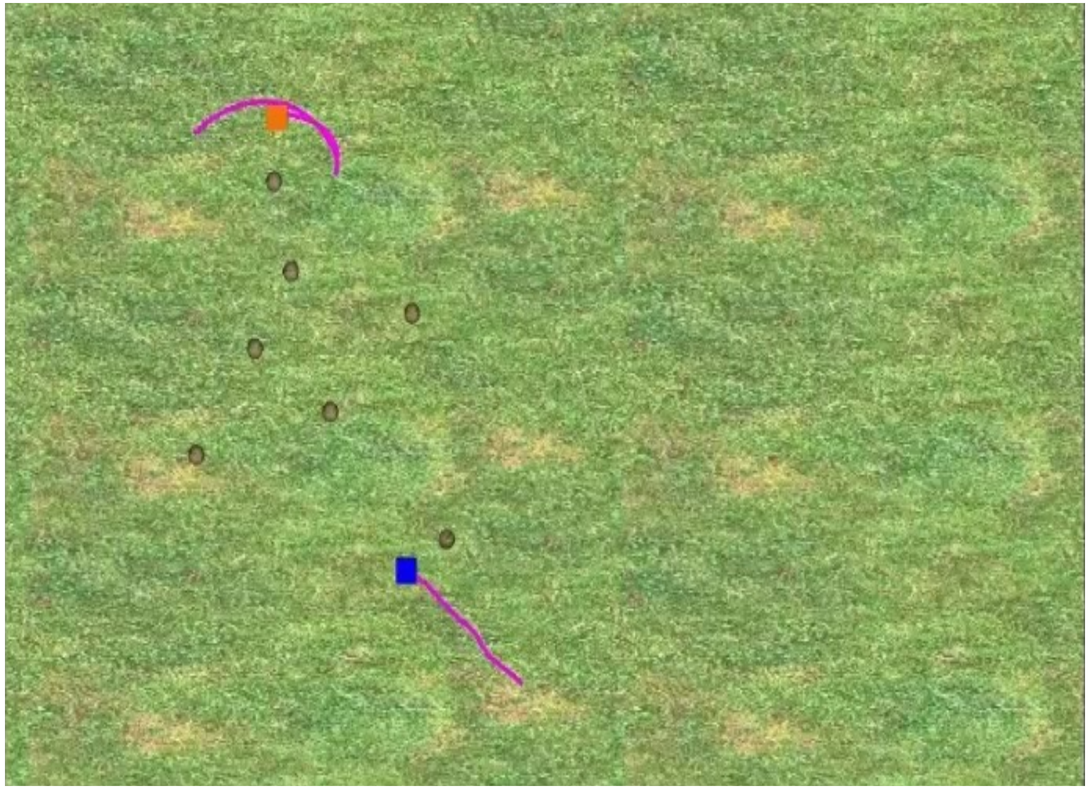

Oscillatory confinement

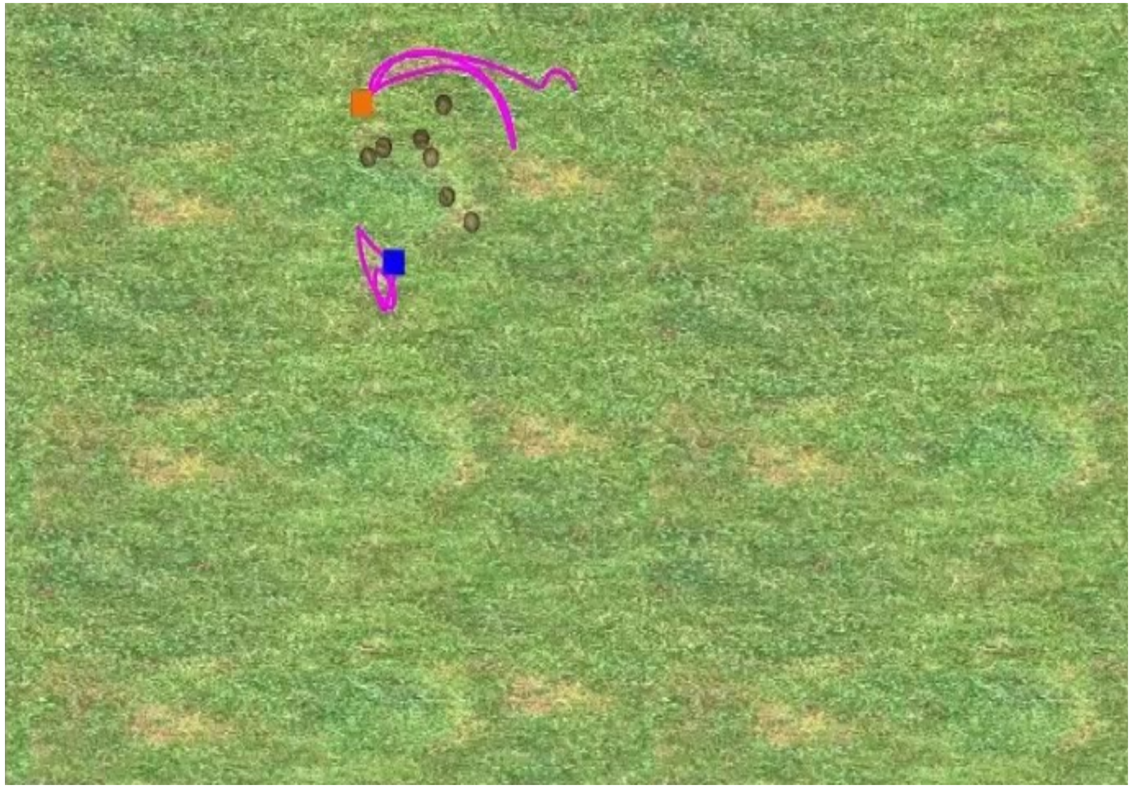

Key limitations of current solutions

- The existing models suffer from one or more of the following limitations:

  1. Ad-hoc modeling assumptions to replicate natural/human behaviour
  2. Assumption of flocking or lack of own dynamics in the targets
  3. Infinite sensing ability of the herders
  4. Lack of scalability as the number of targets increase
  5. Off-line (optimal) computation of the herders behaviour

- More importantly, current solutions do not exploit a key feature of complex systems: the ability of exhibiting emerging behaviour (e.g. oscillatory motion observed in humans)
- Our ongoing work aims at finding solutions to overcome these limitations systematically relaxing these assumptions

Key research questions

Can the emerging collective behaviour of a complex multiagent system (the herders) solve a distributed control task?

Can local simpler feedback rules solve the shepherding problem in the presence of <u>non-cohesive</u> targets?

Under what "herdability" conditions multiple herders can effectively shepherd a group of targets if they only have <u>limited sensing</u>?

Removing cohesiveness from the targets

- We consider *the herders* as distributed feedback control system of the form

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

- To remove any assumption of flocking,
we assume *targets* diffuse randomly in the environment

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

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

- <u>Control goal:</u> $\|\mathbf{x}_i(t) - \mathbf{x}^\star\| \leq r^\star, \qquad \forall i, \forall t \geq t_g,$

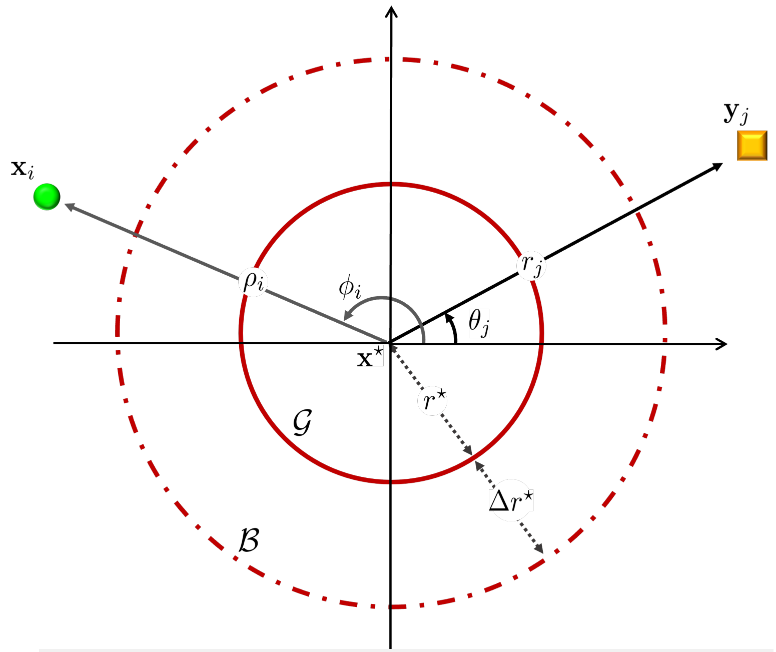

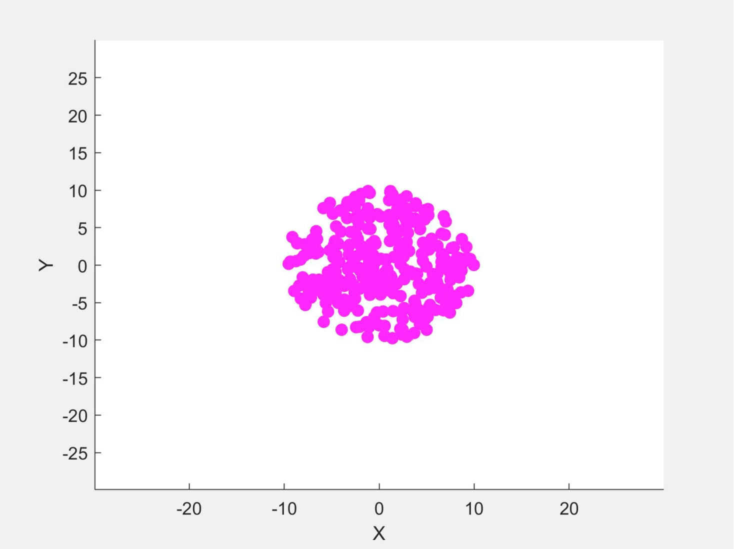

Local control rules

- For the herders we choose (in polar coordinates)

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

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

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

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

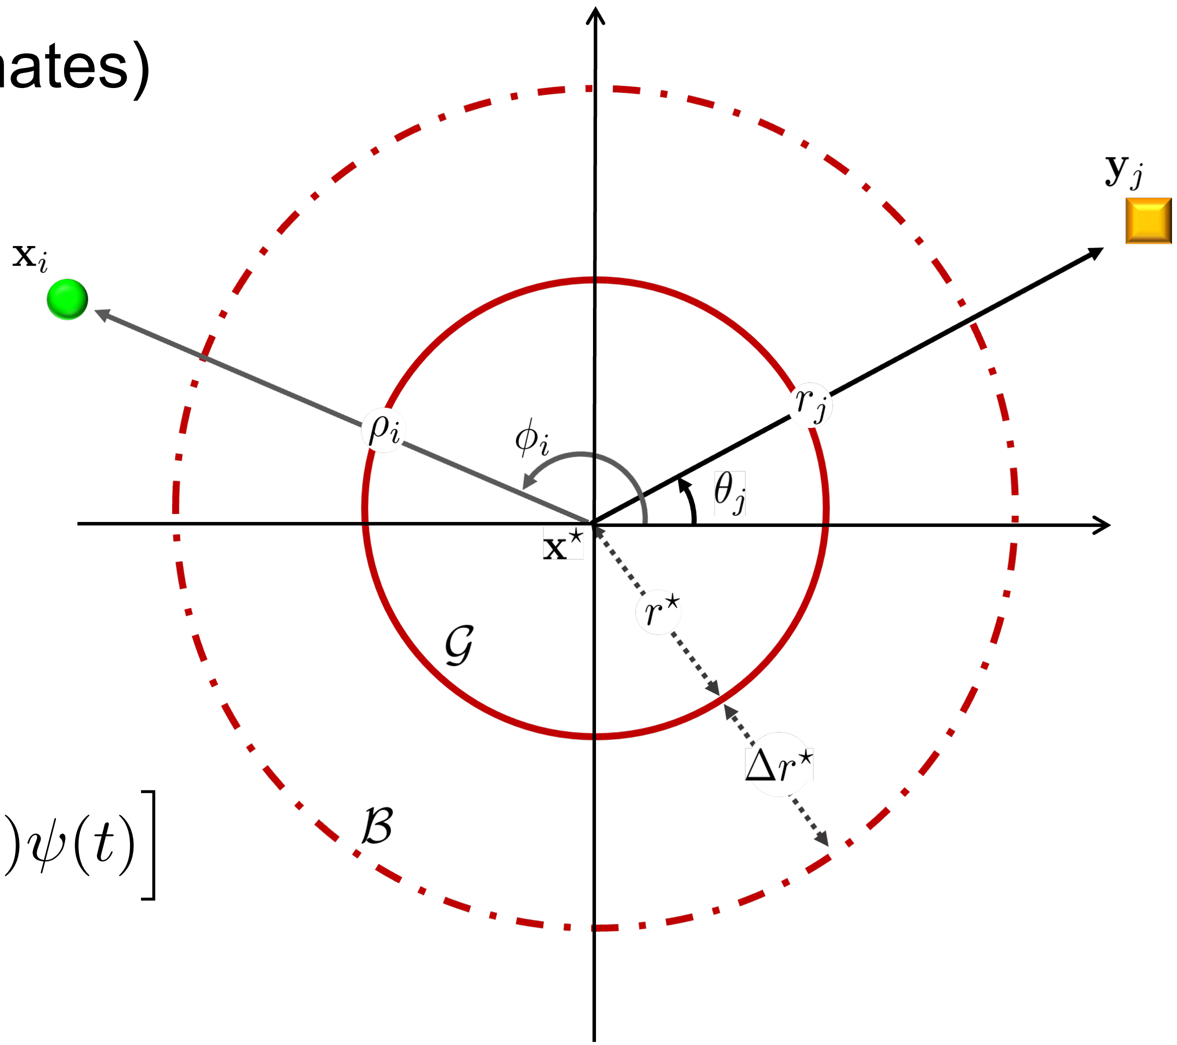

- But this is not enough…

Herder cooperation

- In the presence of multiple herders, it is necessary for herders to negotiate what targets they select (*target-selection*)
- Typically this is done by solving off-line or on-line optmization problems or by using formation control (encirclement etc)
- We want to find simple, yet effective, online **target selection** strategies…
- …allowing herders to cooperatively select their targets..
- .. without requiring any computationally expensive solution of on-line optimization problems

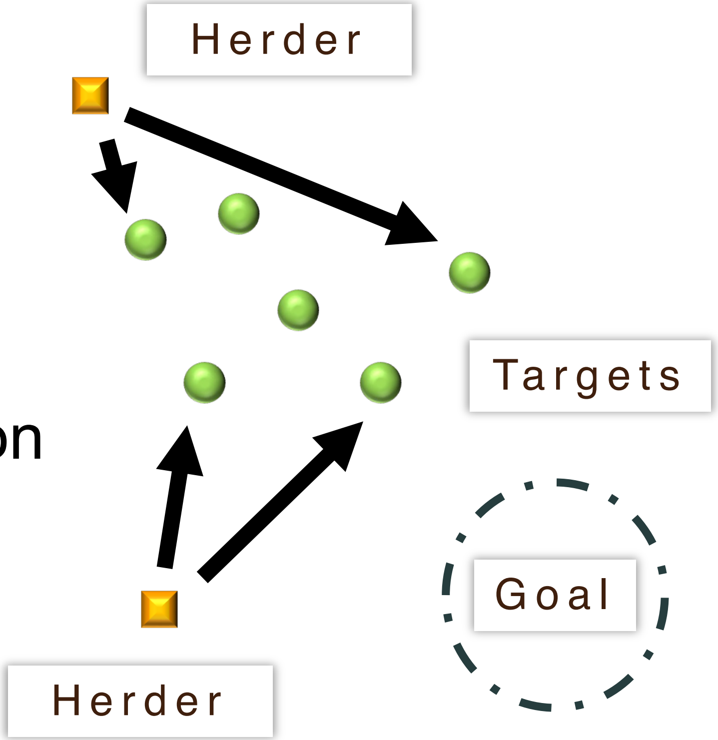

Simplest static strategies

Global search

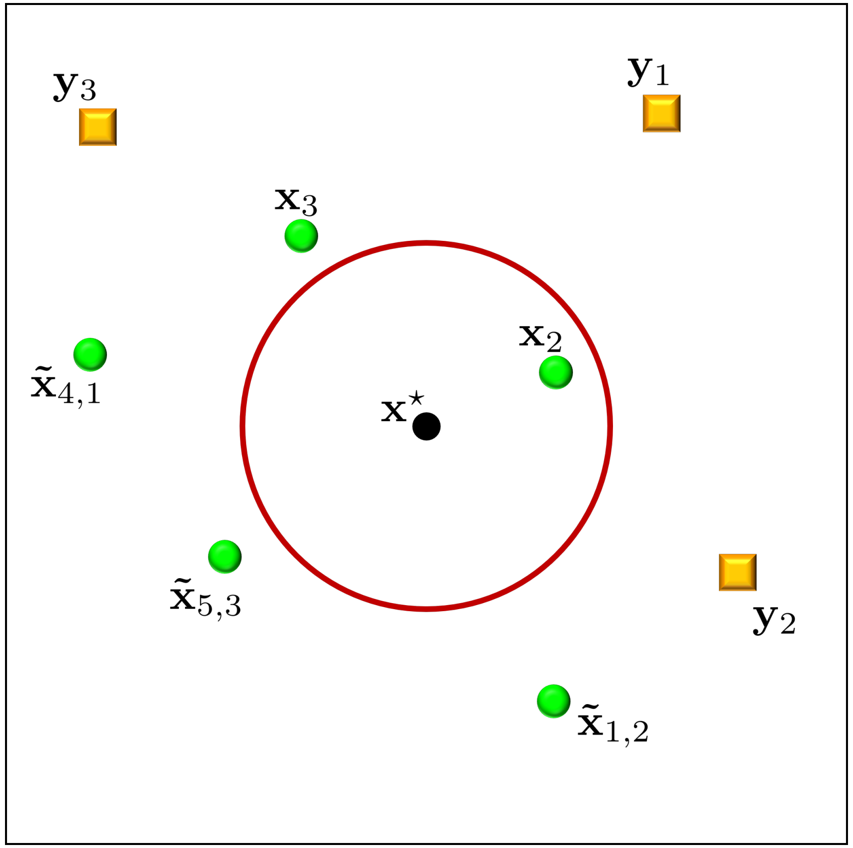

Static arena partitioning

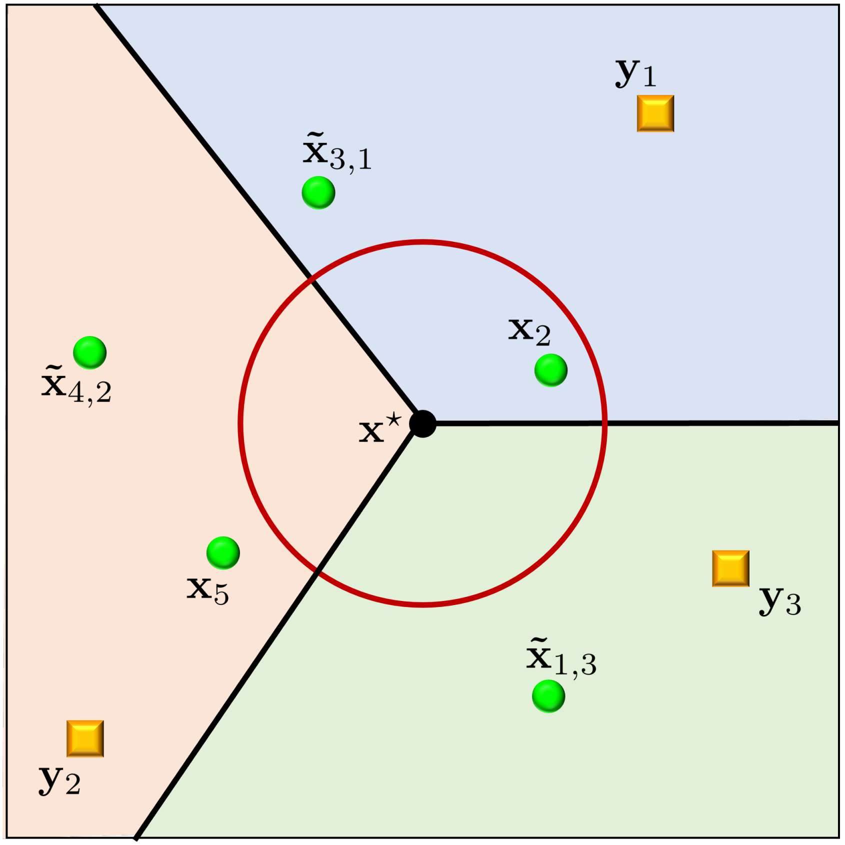

Dynamic selection strategies

- *Leader-follower target selection strategy*
- At start, herders are labelled anticlockwise starting from a randomly selected herder (the leader)
- The plane is then partitioned dynamically in different search regions of **constant width** for each herder

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

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

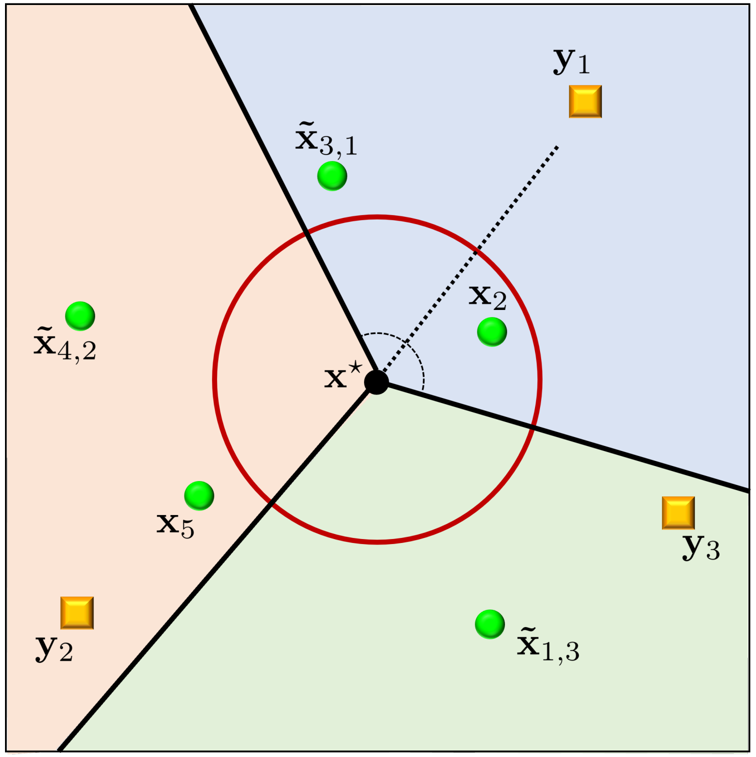

- As the leader chases a target, the other herders' regions adjust dynamically

Peer-to-peer target selection strategy

- In this case the **width** of the sectors assigned to each herder is also **dynamically changing**

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

  as a function of the relative angular distance between neighboring herders

- Note that in this case the herders can self-determine their circular sector of interest by just observing the relative positions of their neighbors

- Hence they dynamically cooperate to decide who herds whom!

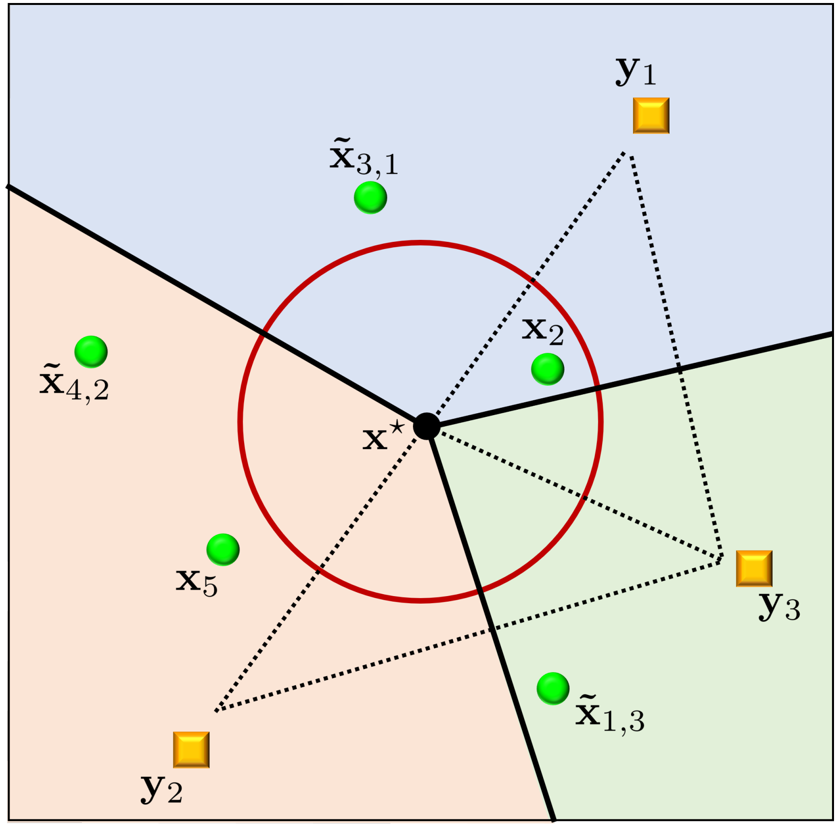

The resulting herding strategy

- Our strategy consists therefore of local control laws driving the dynamics of each herder and a target selection strategy allowing them to somehow cooperate

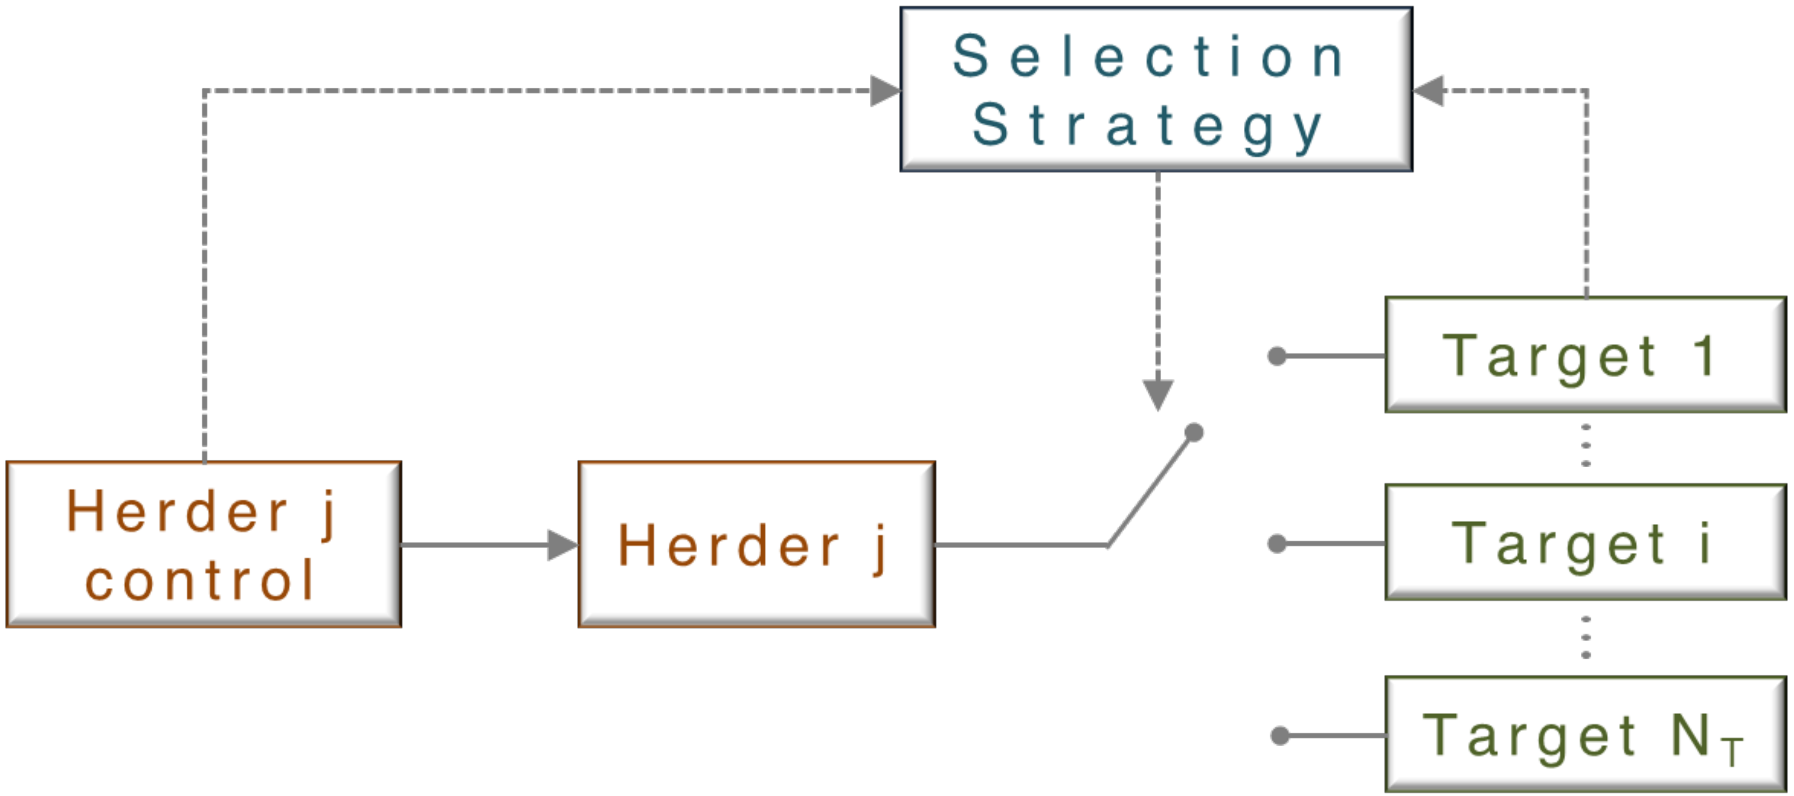

Experimental validation

- We validated our herding strategy via both numerical simulation and experiments

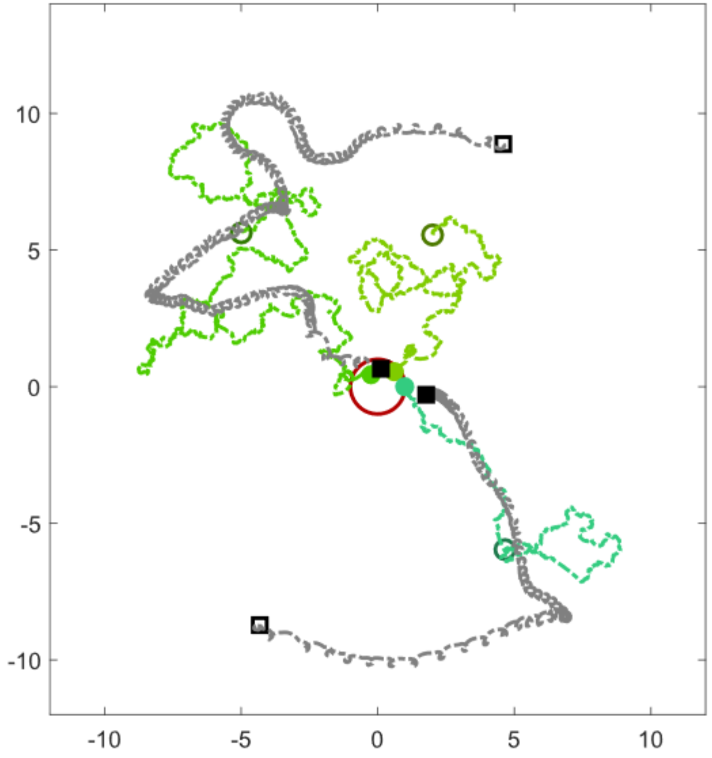

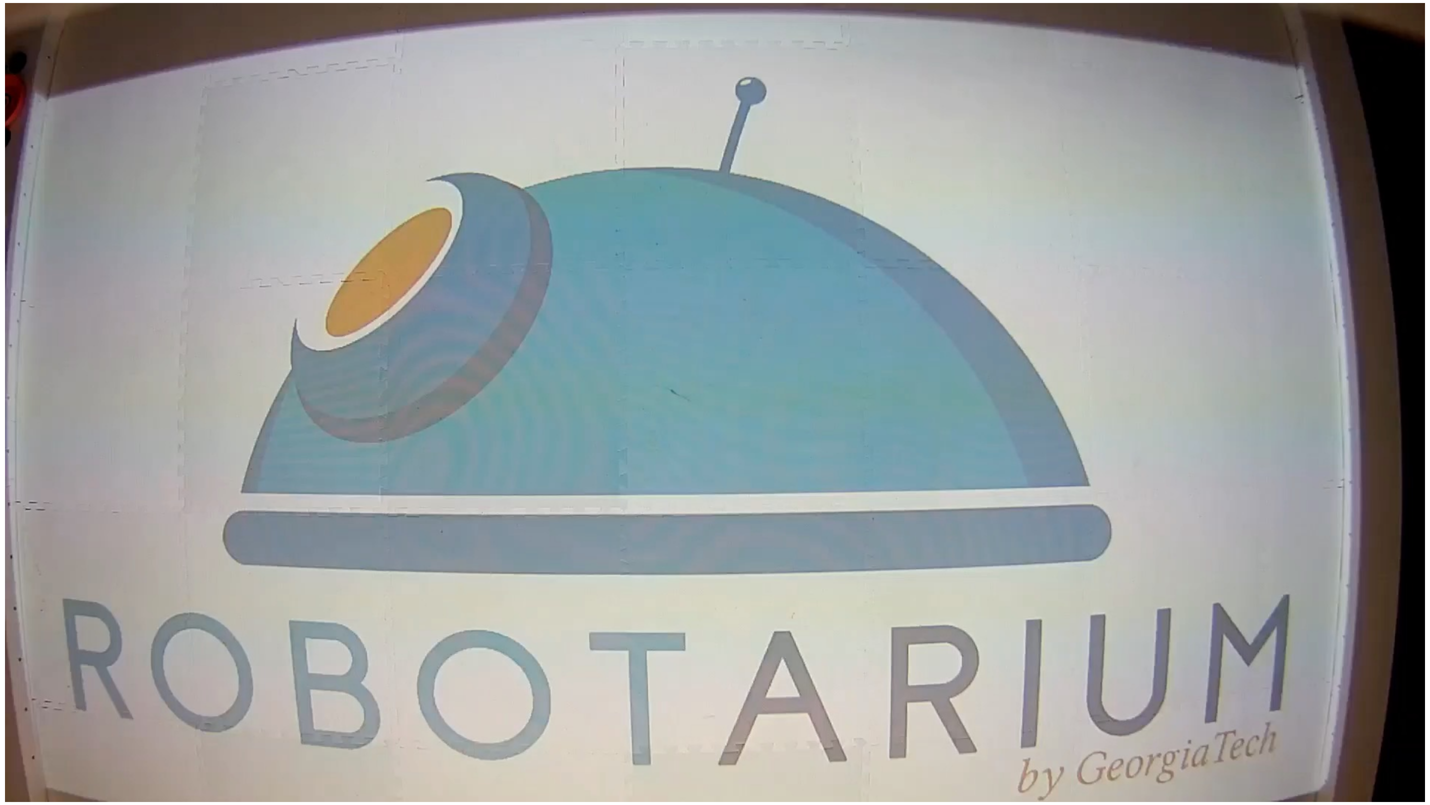

Sean Wilson, et al., ["The Robotarium […]"]() in IEEE Control Systems Magazine, vol. 40, no. 1, pp. 26-44, Feb. 2020.

Overview of the validation results

	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

average length d_g of the path travelled by the herders

gathering time t_g

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

the herd agents' spread $S_{\%}$

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

the mean distance D_T

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

	S&R	COC
Global	100%	0%
Static	0%	100%
Leader-follower	12%	88%
Peer-to-peer	68%	32%

Scalability and Robustness

- Containment time

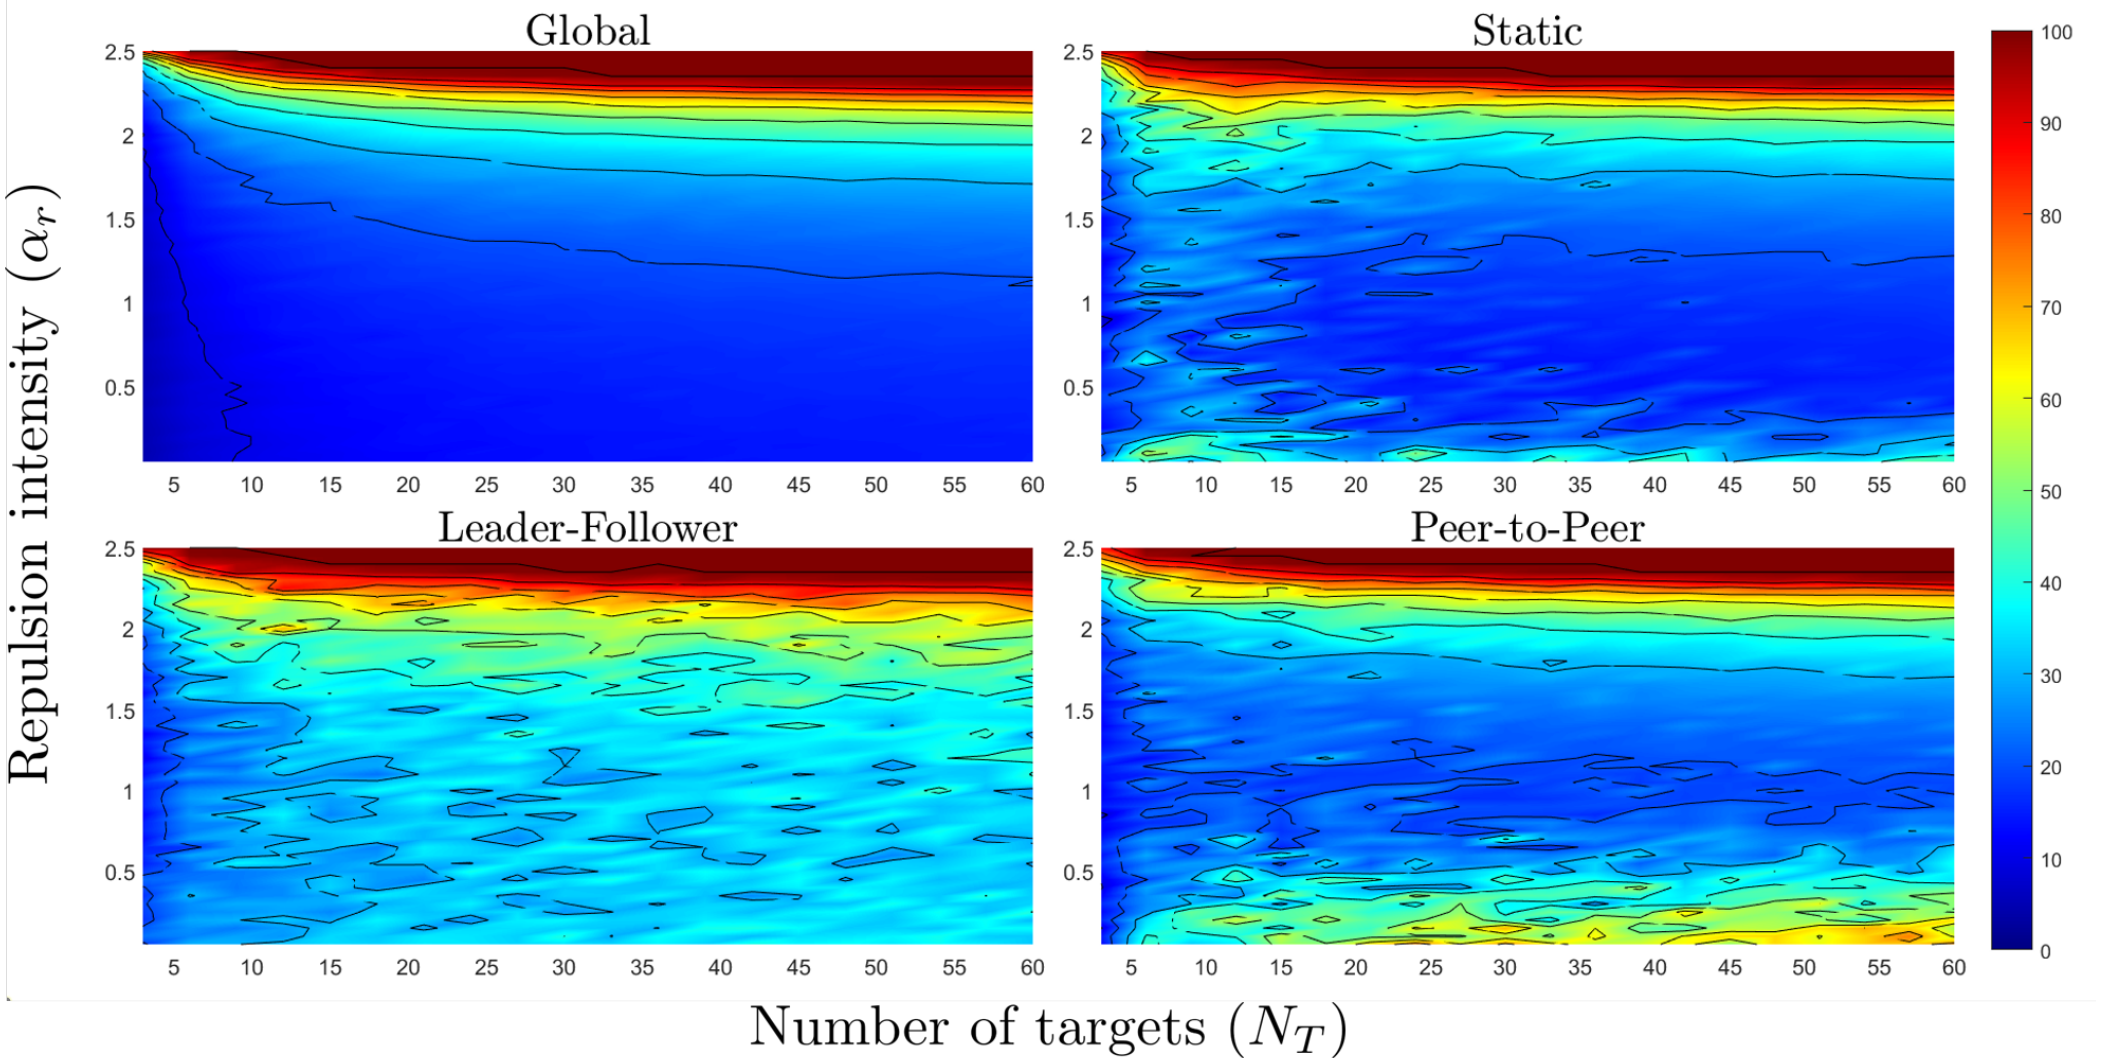

To summarize

- We were able to solve the herding problem via a set of simpler local rules driving the individual herders' behaviour..
- ..complemented by target selection rules
- <u>Still *we assumed **global** rather than limited sensing* of the herders</u>
- What if the herders only possess limited sensing?
- How many targets can they shepherd?

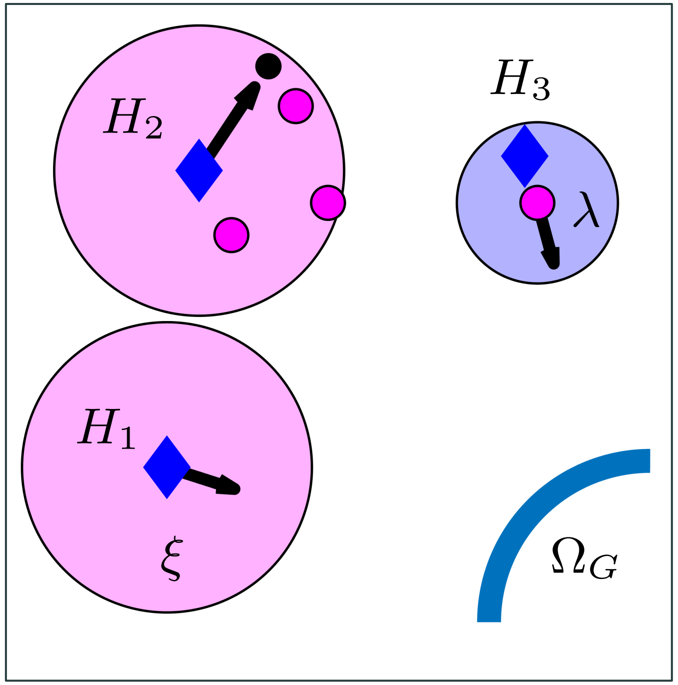

A minimal model of shepherding with limited sensing

- *M* targets, *N* herders initially randomly distributed

$$\dot{\mathbf{T}}_a = \sqrt{2D}\boldsymbol{\mathcal{N}}$$

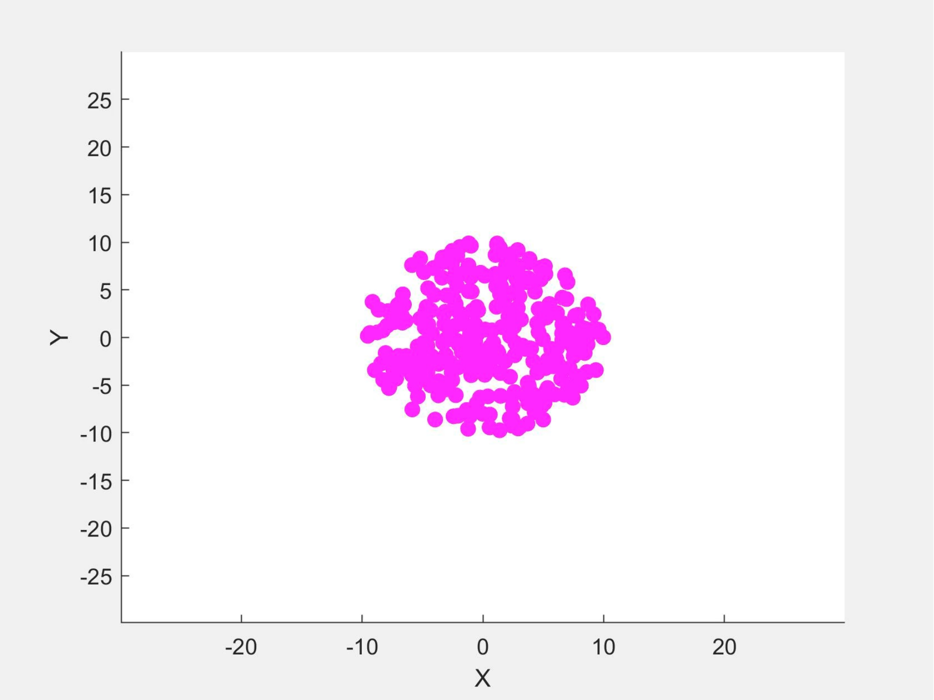

$$\beta\lambda^2 \gg D$$

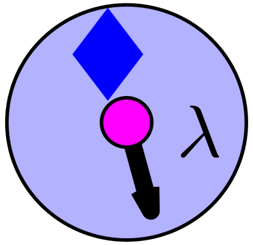

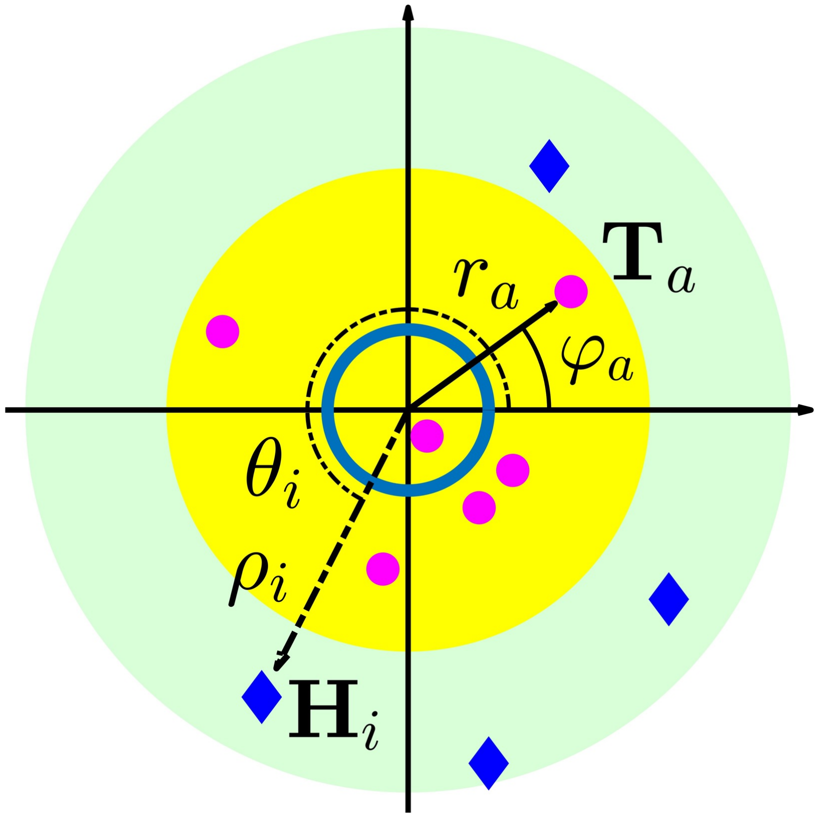

A. Lama, MdB, "Shepherding control and Herdability in Complex Multiagent Systems", *Physical Review Research,* 2024

Herders' local dynamics

$$\dot{\mathbf{H}}_i = (1-\eta_i)\mathbf{F}_i(\mathbf{H}_i, \mathbf{r}^*) + \eta_i \mathbf{I}_i(\mathbf{T}, \mathbf{H}, \xi)$$

$$\mathbf{F}_i(\mathbf{H}_i, r^*) = \begin{cases} -v_H \widehat{\mathbf{H}}_i & \text{if } |\mathbf{H}_i| \geq r^* \\ 0 & \text{otherwise} \end{cases}$$

$$\mathbf{I}_i(\mathbf{T}, \mathbf{H}, \xi) = -\left[\alpha\left(\mathbf{H}_i - (\mathbf{T}_i^* + \delta\widehat{\mathbf{T}}_i^*)\right)\right]_{v_H}$$

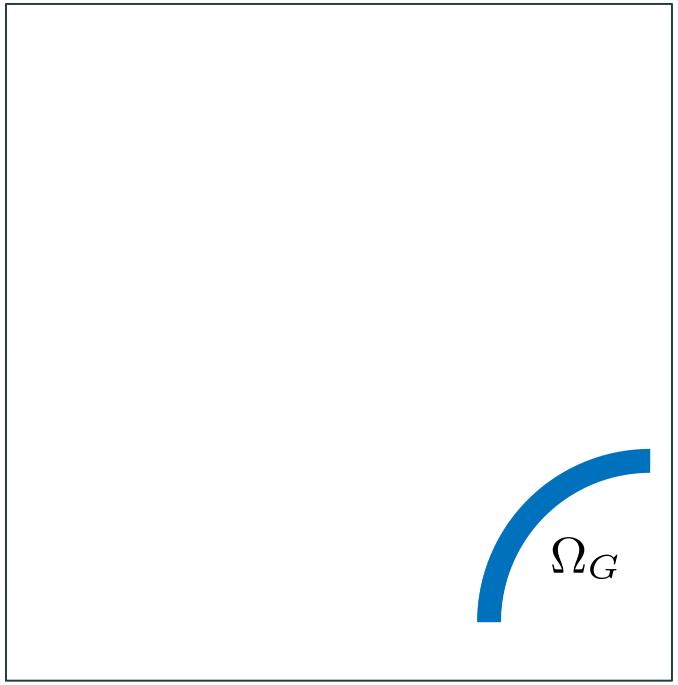

- At every step, herders *make decisions* on what target to chase
- Reciprocal interactions can also be added, e.g. collision avoidance

Target selection rule

- The target to chase is selected by an herder as the target with the largest distance from the origin within its sensing region
- If an herder detects other herders in its sensing region

$$|\mathbf{H}_j - \mathbf{H}_i| \leq \xi$$

  it will only considers those targets such that

$$|\mathbf{T}_a - \mathbf{H}_i| \leq |\mathbf{T}_a - \mathbf{H}_j|$$

- Nearby herders effectively cooperate through this simple rule

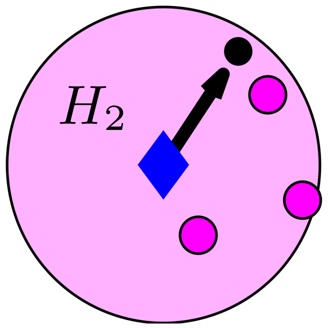

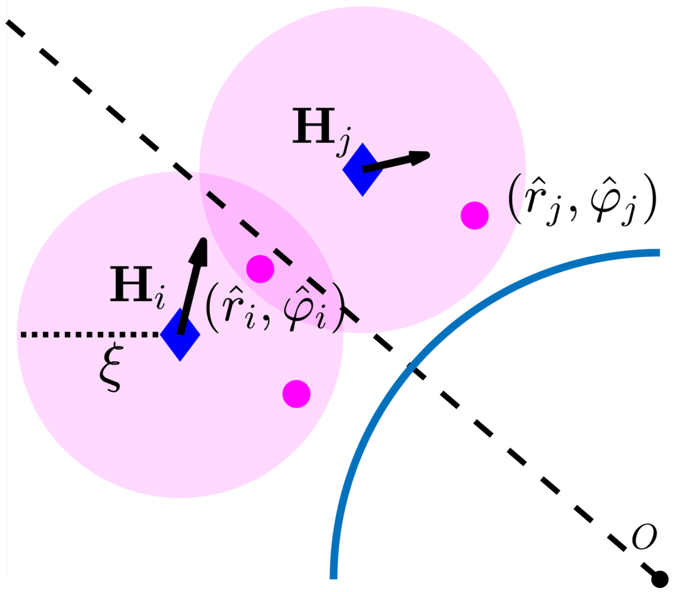

Shepherding can be successful

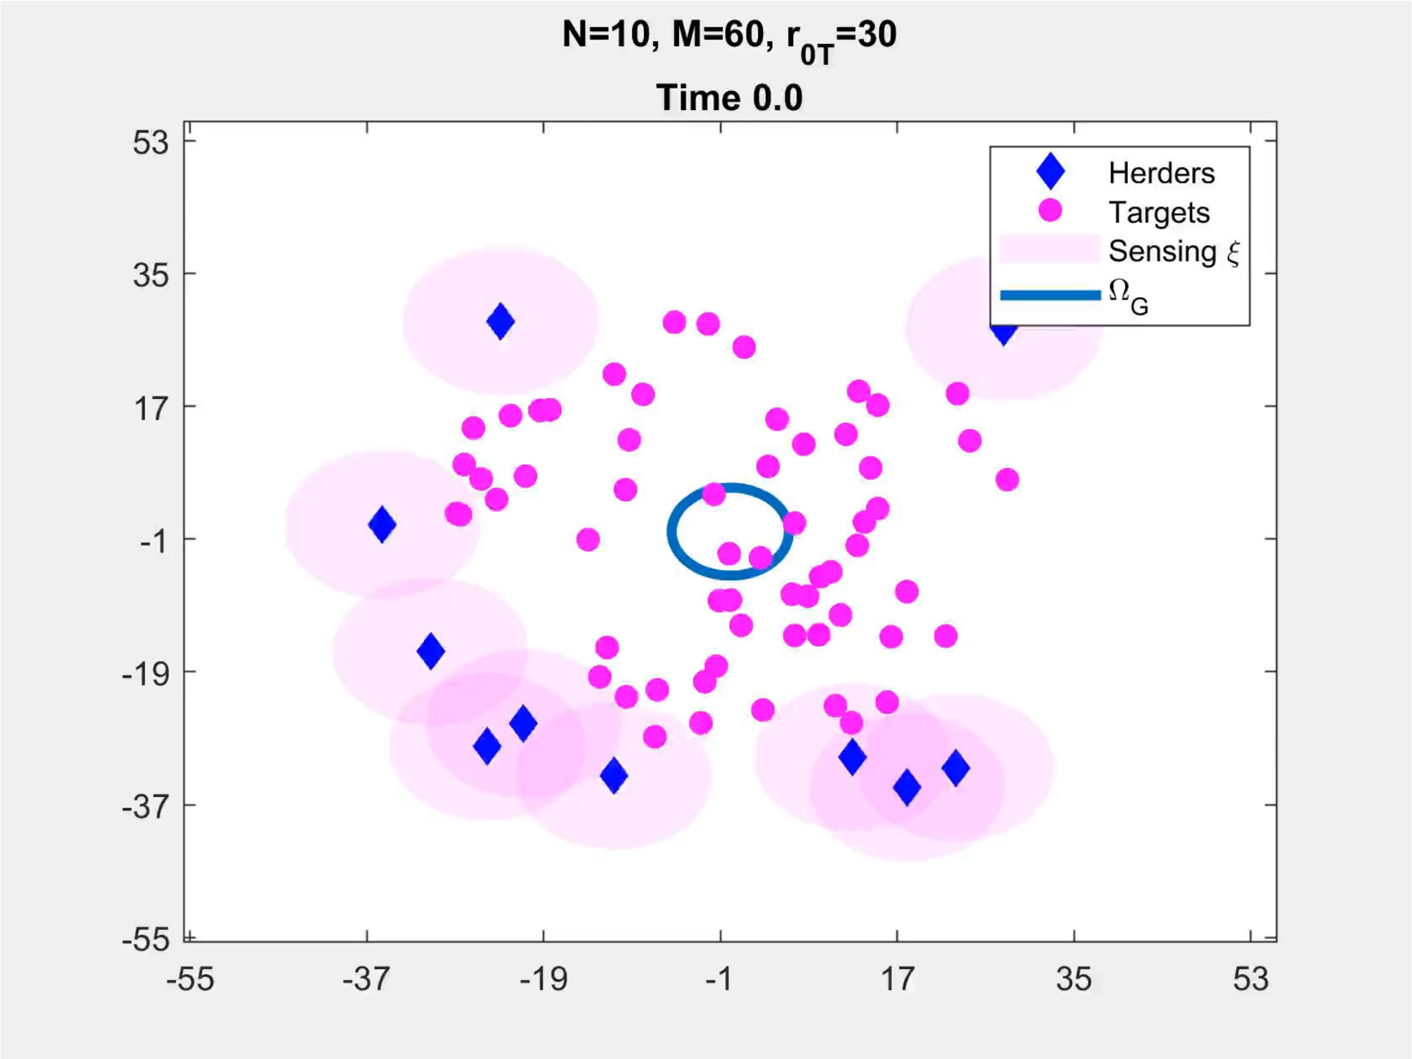

$$v_H > v_T$$

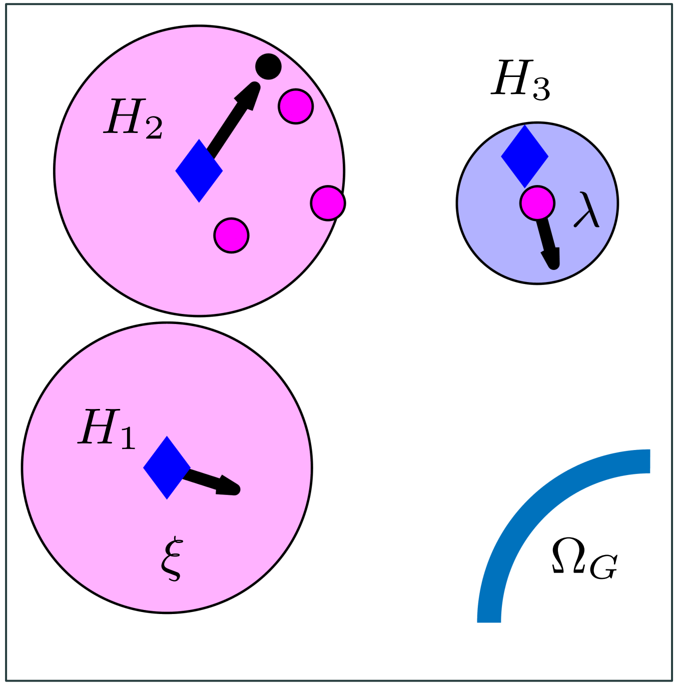

The herdability problem

- Under what conditions on the repulsion zone, the sensing area and the density of the targets can we achieve herdability of a given number of targets?

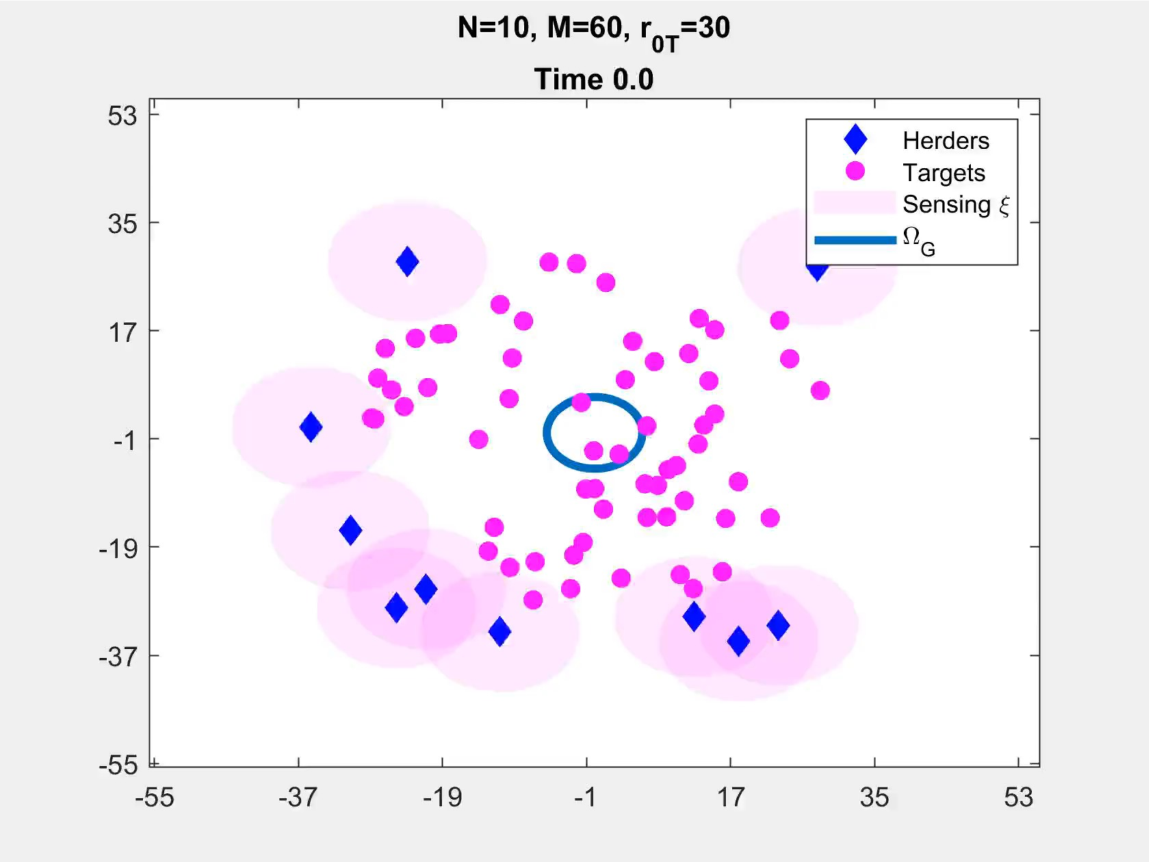

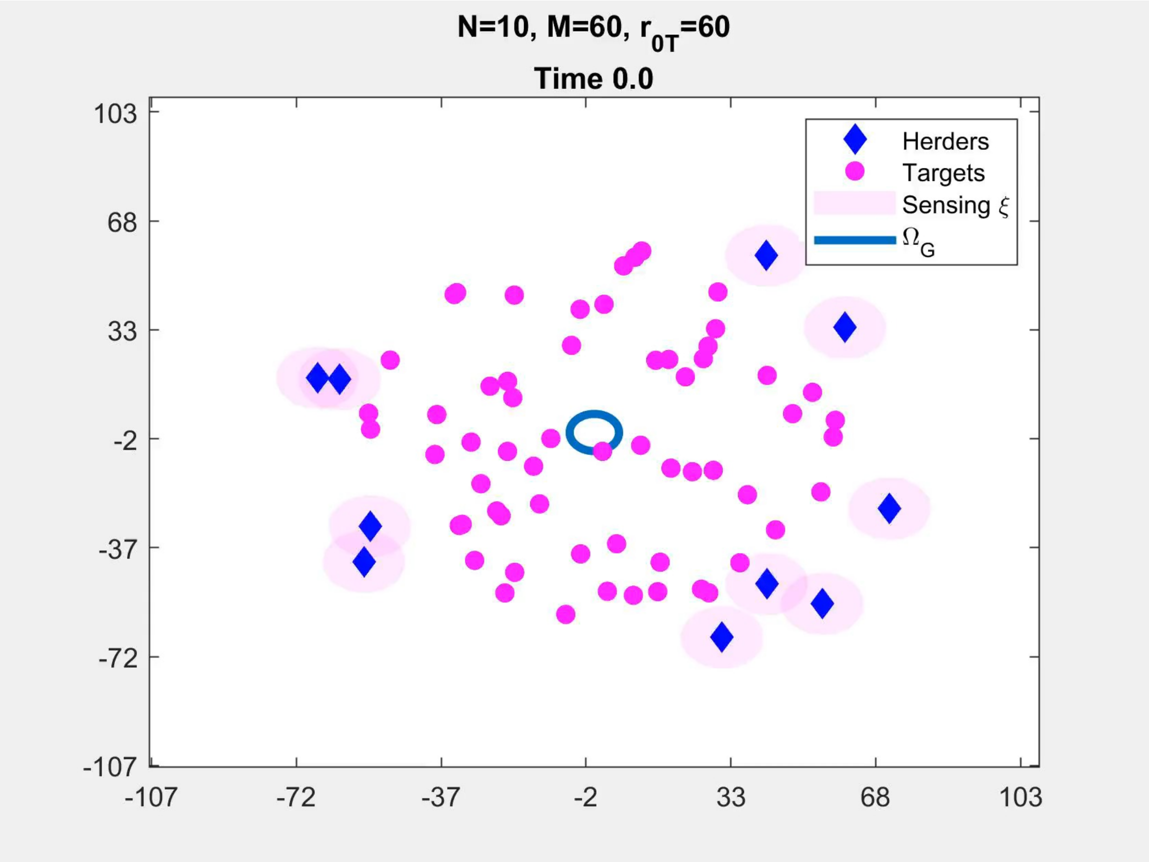

A. Lama, MdB, ”Shepherding control and Herdability in Complex Multiagent Systems”, *Physical Review Research,* to appear

Herdability

- We define a group of *M* target agents as "herdable" by *N* herders if the latter can guide a significant fraction of the former within a finite time towards the goal region
- We look for the minimum number of herders $N^*(M)$ necessary to herd M targets
- The *herdability chart* for <u>infinite sensing</u> shows: $N^*(M) \propto \sqrt{M}$

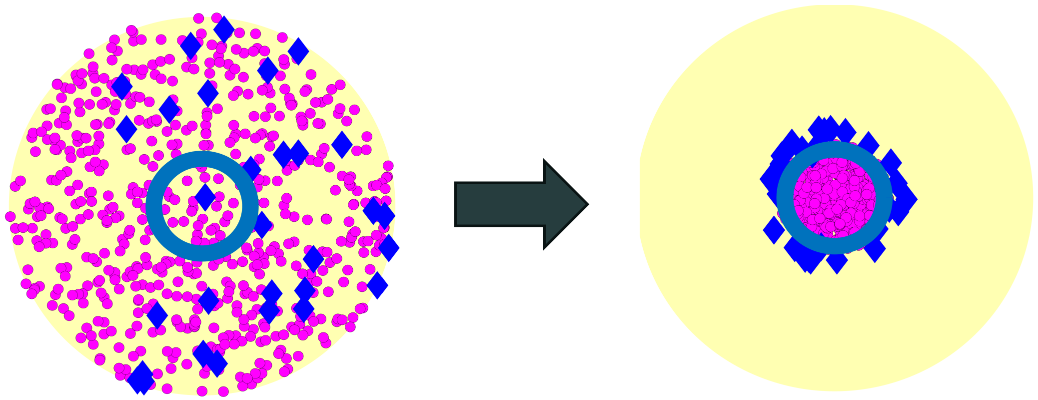

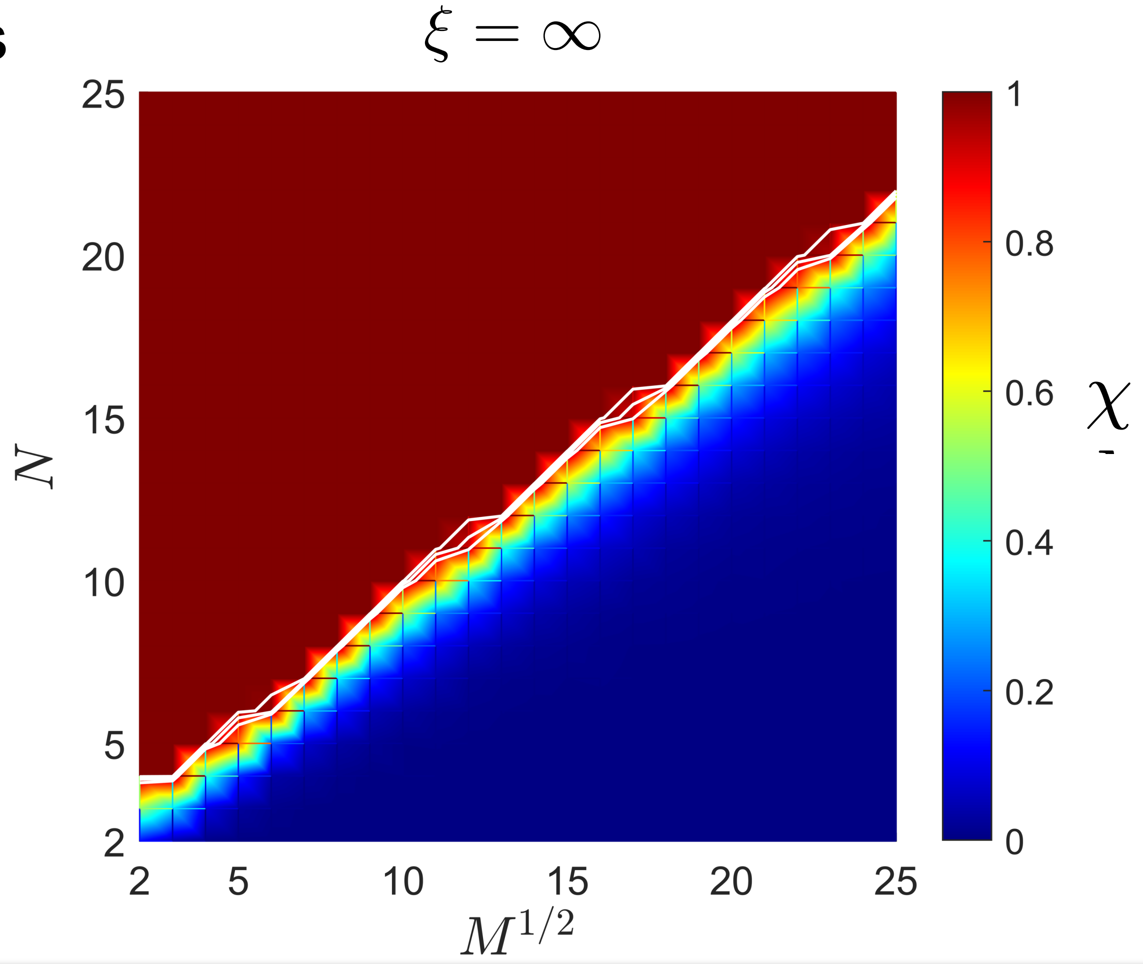

Finite sensing effect

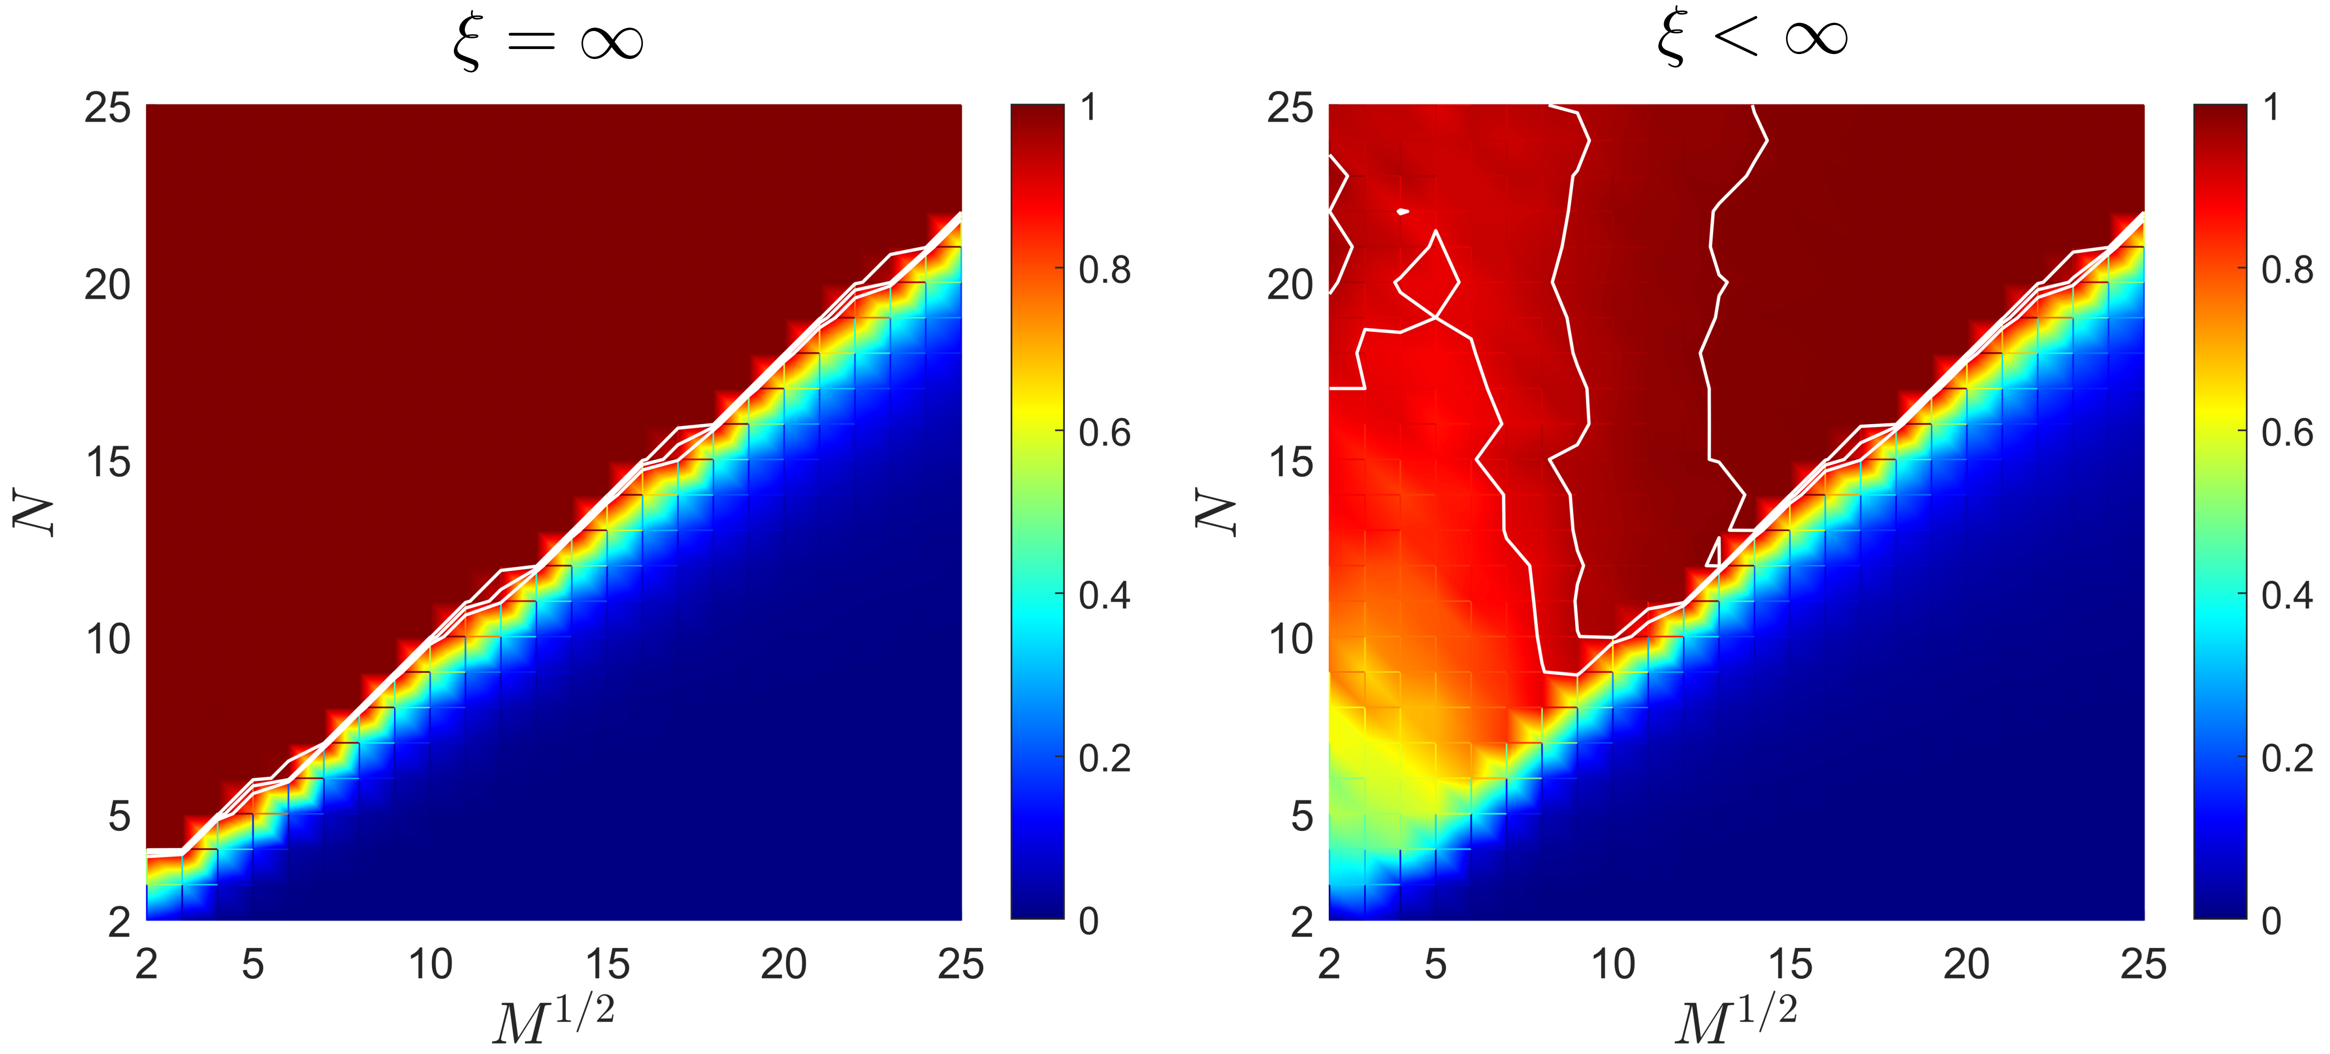

Fewer targets do not necessarily ease the shepherding task!

Herdability conditions

- In general, in order to be successful herders need to:

  1. *Collectively sense all the targets* which are random independent walkers

  2. *Counterbalance diffusion of the M targets* with the transport flow they induce

- In *the finite sensing case*, meeting the first condition becomes harder and harder as targets' density decreases and targets become more sparse
- To explain the critical threshold we need to analyze how herders, using only local information, can sense and corrall also distant targets

The herdability graph

- We define the herdability graph as the *random geometric graph*

$$G_{ab}(\boldsymbol{T}, \xi) = 1$$

$$\text{if } |T_a - T_b| \leq \xi$$

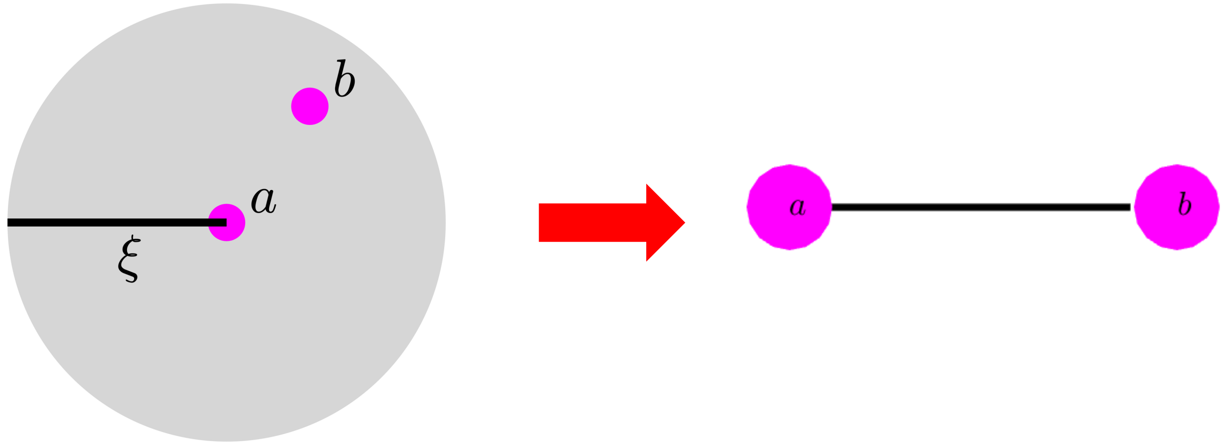

- A path in the graph indicates the potential for a herder to transition from sensing one target to sensing another

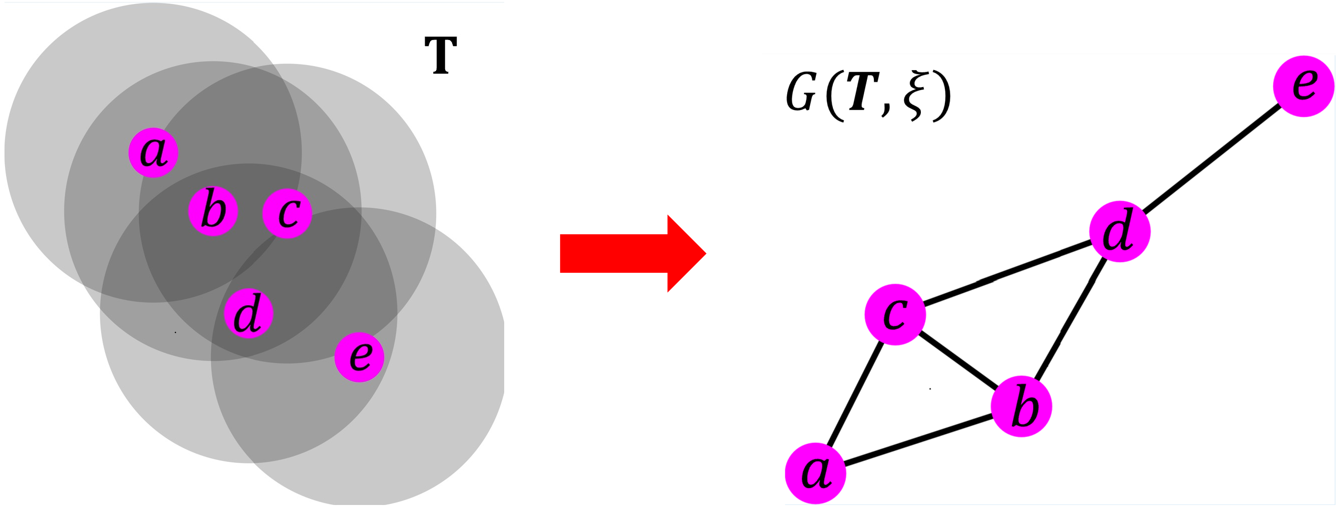

Percolation of the herdability graph

- Namely if the graph is too sparse, herders cannot navigate to reach the furthermost targets
- Hence, some targets will be lost
- Hence, we proposed to estimate the critical threshold by studying *percolation of the herdability graph*

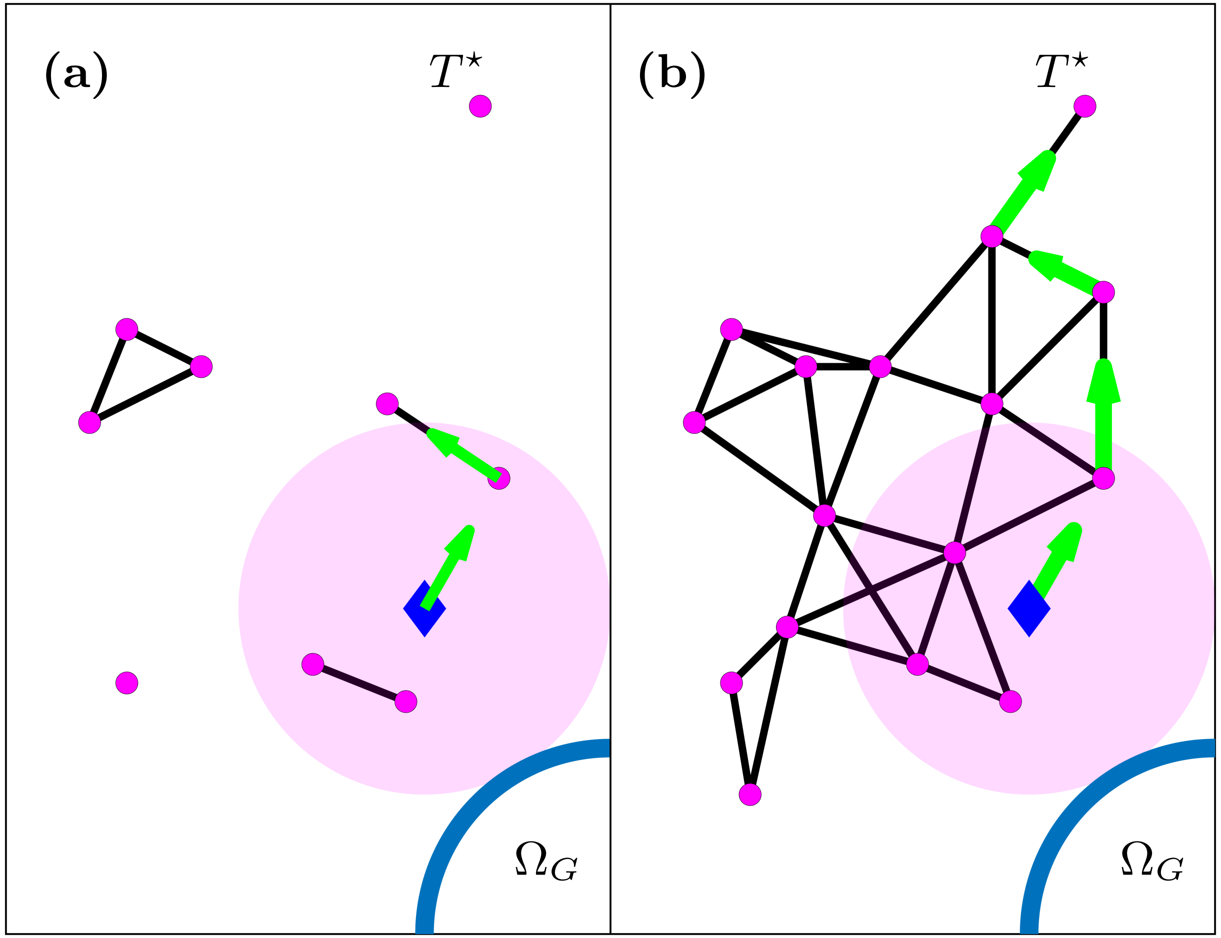

Percolation analysis

- We study when G becomes initially connected when targets are randomly distributed in a circle of radius R
- We then expect

$$M^{\mathrm{low}} \sim R^2/\xi^2.$$

- This aligns with our numerical findings explaining the observations and the scaling observed in the numerical experiments

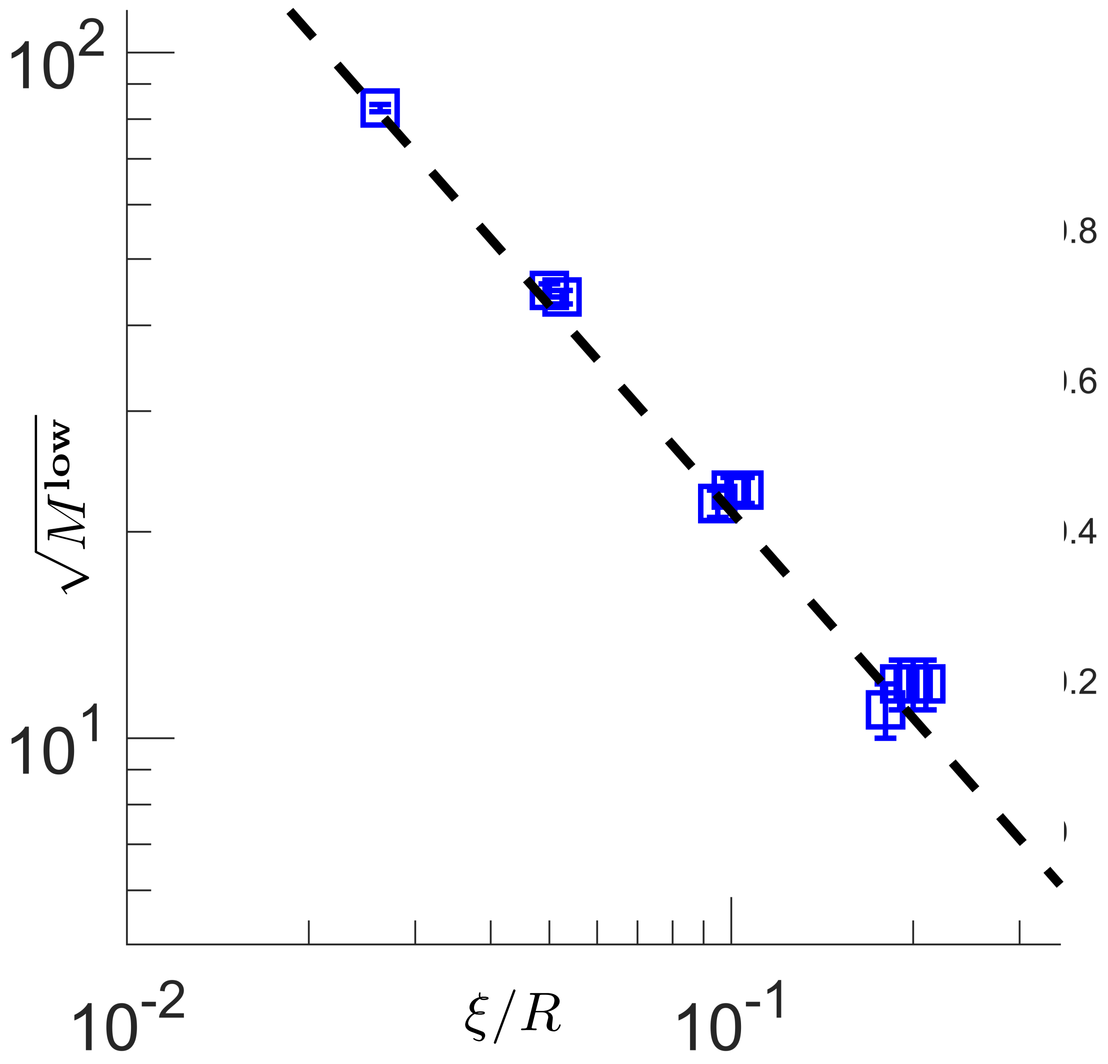

To summarize

- In the presence of limited sensing, some herdability conditions must be satisfied
- As the number of targets increases the problem becomes harder and harder
- More herders are needed to solve the problem
- How dependent is this from the specific rules and dynamics we selected?
- Can we find other (better) solutions to the problem?
- To address this issue we started exploring the use of learning-based approaches

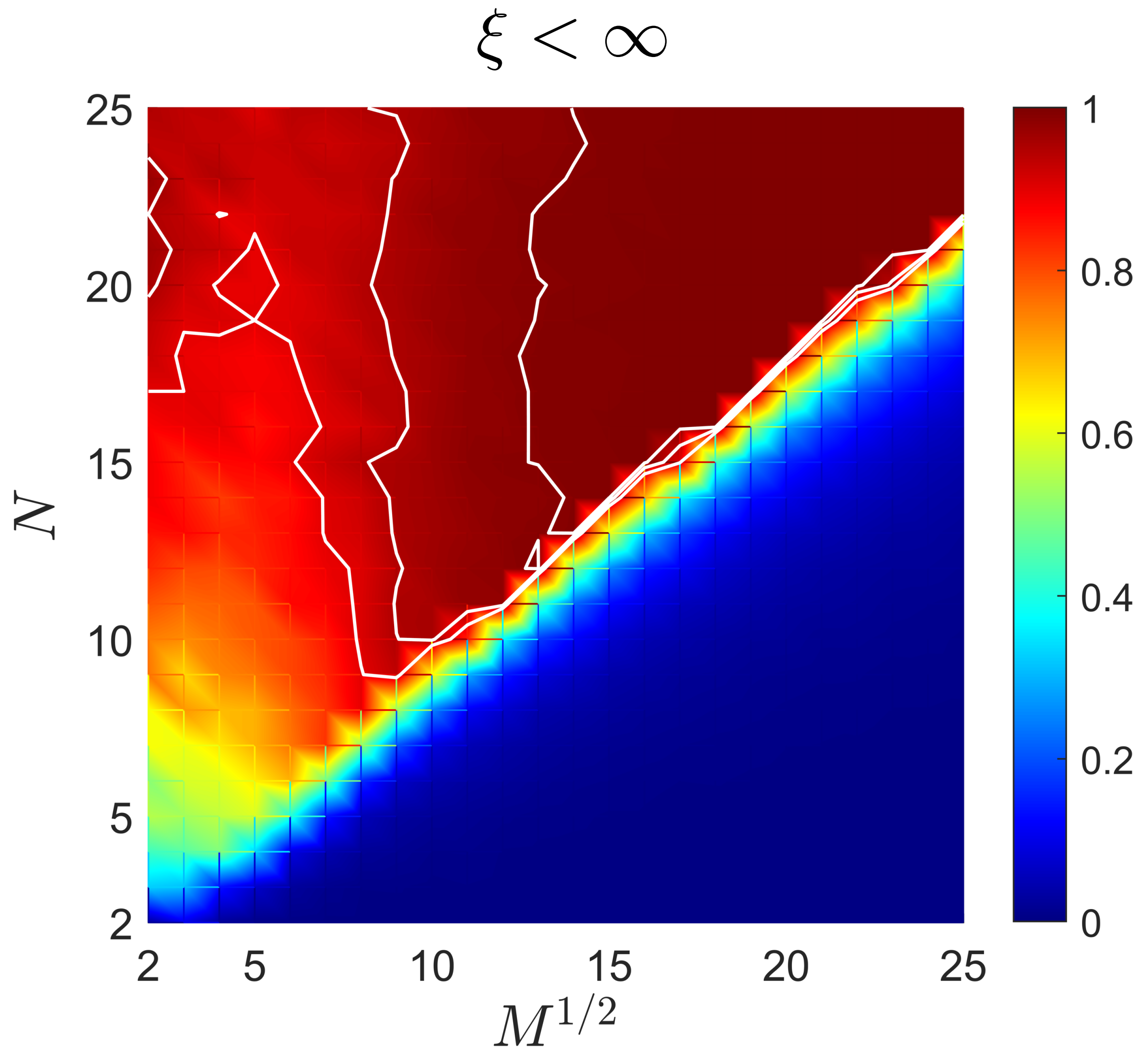

A computational approach

- Shepherding is a ”multi-layer” control problem consisting of three layers *(coralling, target selection and herder cooperation)*

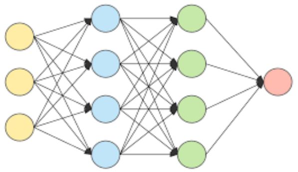

	Task		Scenario
L1	Steering and containing a target	⟺	Single herder single target
L2	Selecting the target to steer	⟺	Single herder multiple targets
L3	Coordinating the herders	⟺	Multiple herders multiple targets

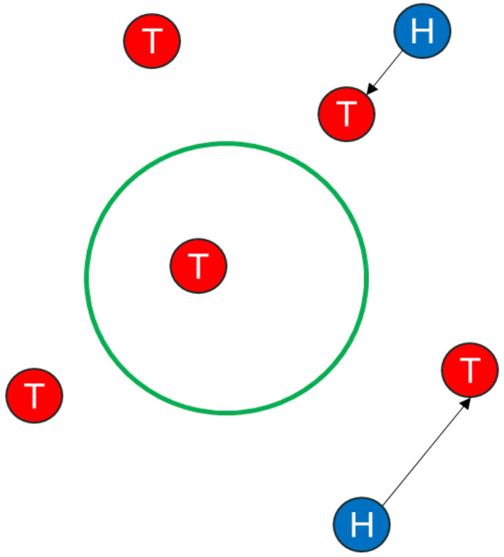

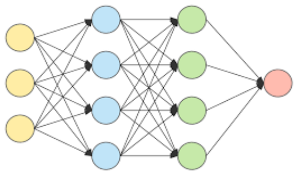

Two RL driven cooperating herders

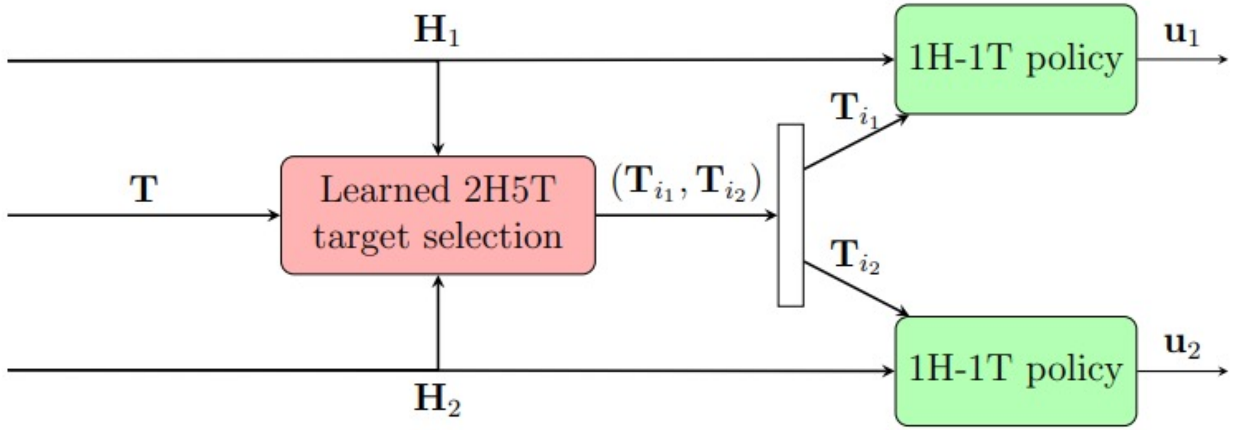

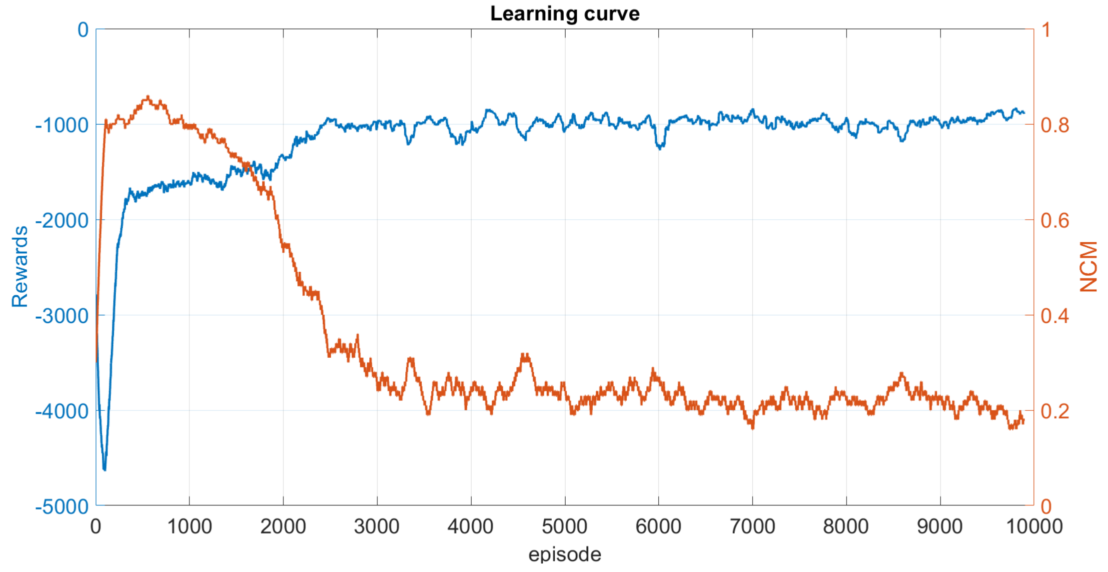

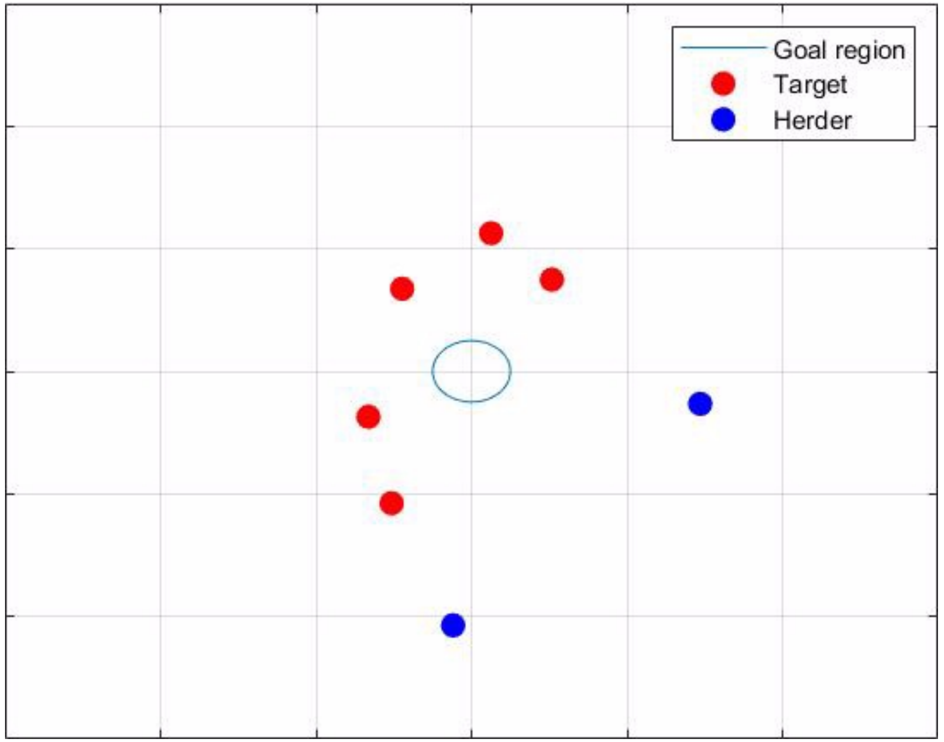

Beyond two herders: the multiple herders case

- DQN has fixed structure: 2H-5T net previously trained is extended (validation only)
  - Each herder only considers the 5 closest targets and the closest herder (**topological sensing**)
- Unbounded region (with I.C. in circle of radius 30m)

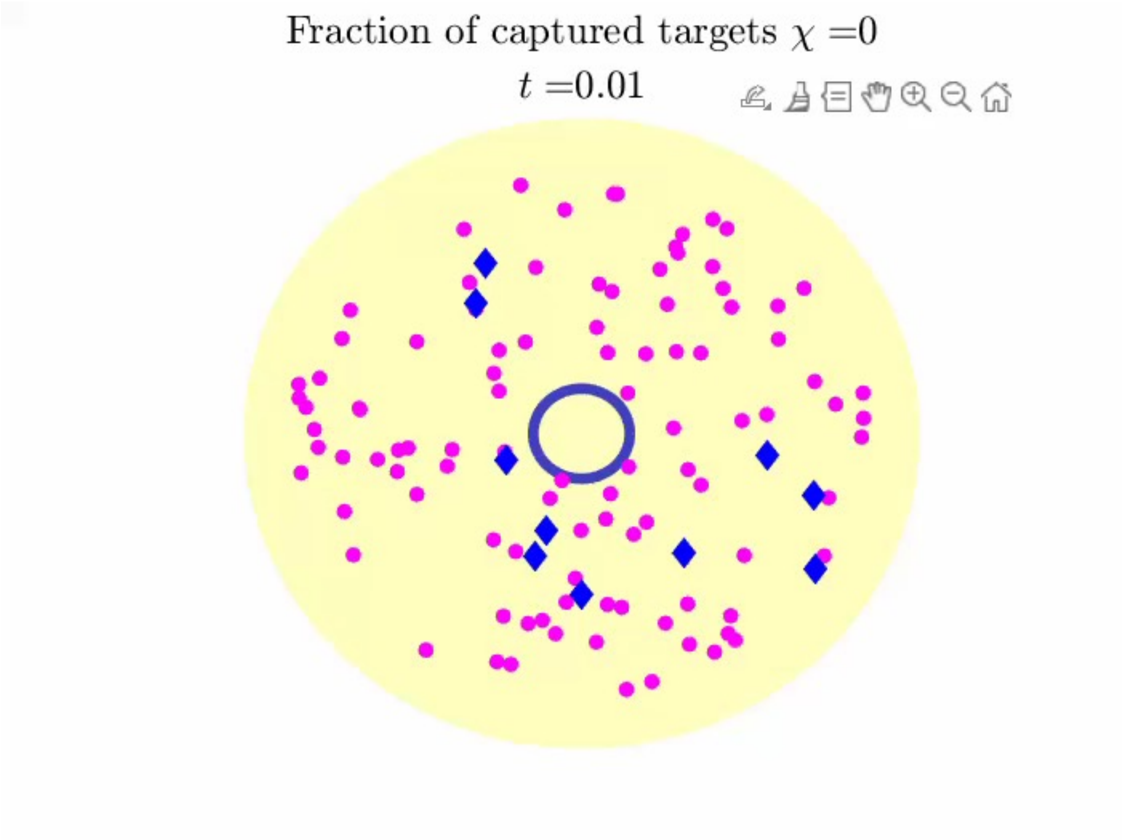

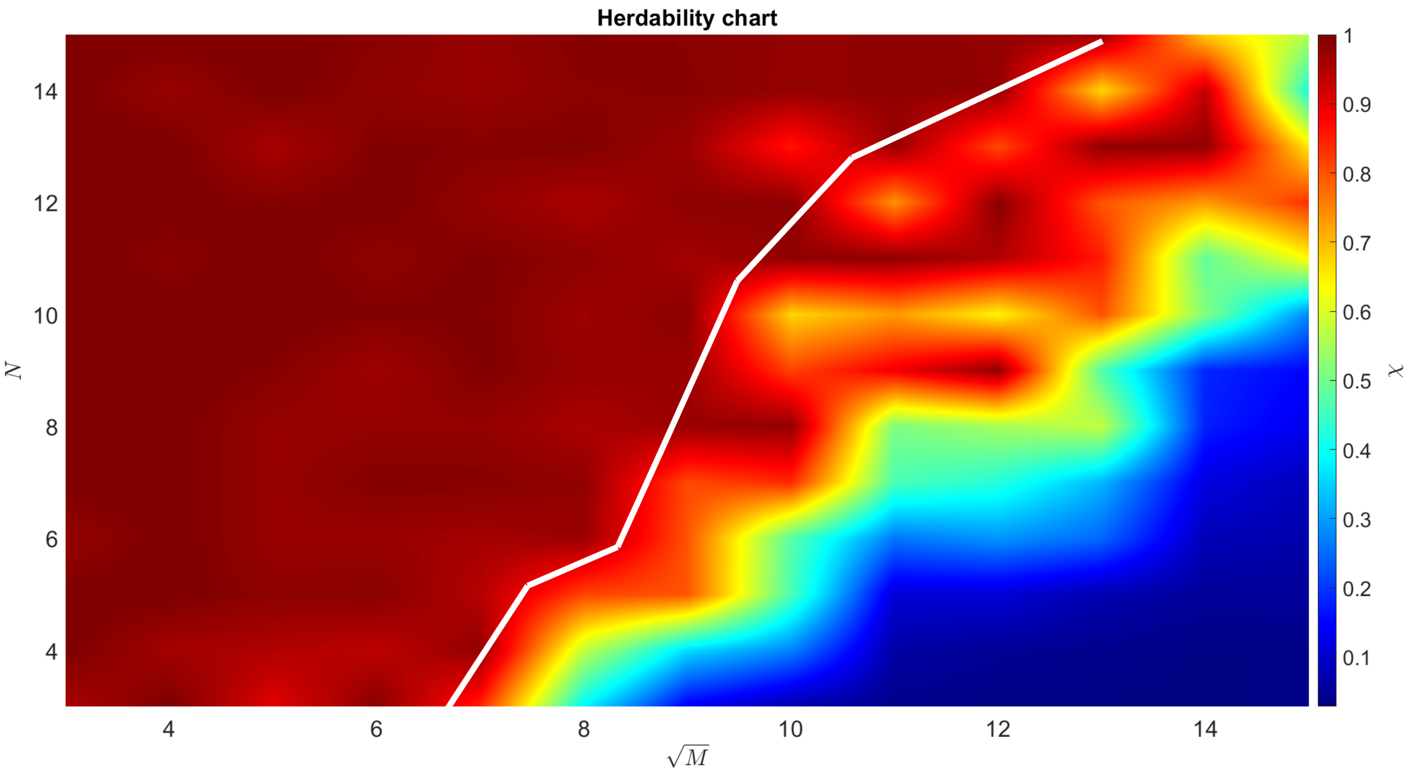

Conclusions

- Complex systems can be exploited to solve distributed control tasks
- *Shepherding* is a great paradigmatic example..
- ..where emerging behaviour needs to arise from a complex system in order to solve a control task
- We saw that a *set of simple local rules* coupled with appropriate *target selection strategies* is able to solve the problem under certain assumptions

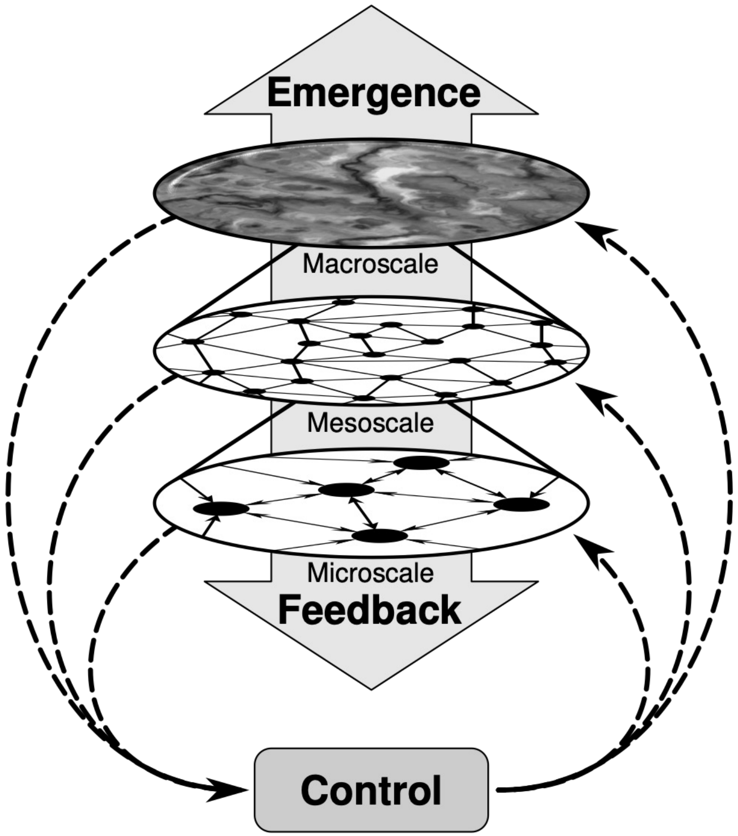

- We studied the herdability of the targets when the *herders' sensing is finite* and linked it to the percolation of a suitably defined *herdability graph*
- Finally, we discussed possible learning-based approaches to solve the problem

Perspectives and open problems

- Many aspects of this problem remain unsolved
- How to prove convergence of a herding policy in the presence of limited sensing?
- Is it possible to give herdability conditions for more general dynamics?
- Can feedback rules be defined in the continuum approximation, e.g. PDE control?

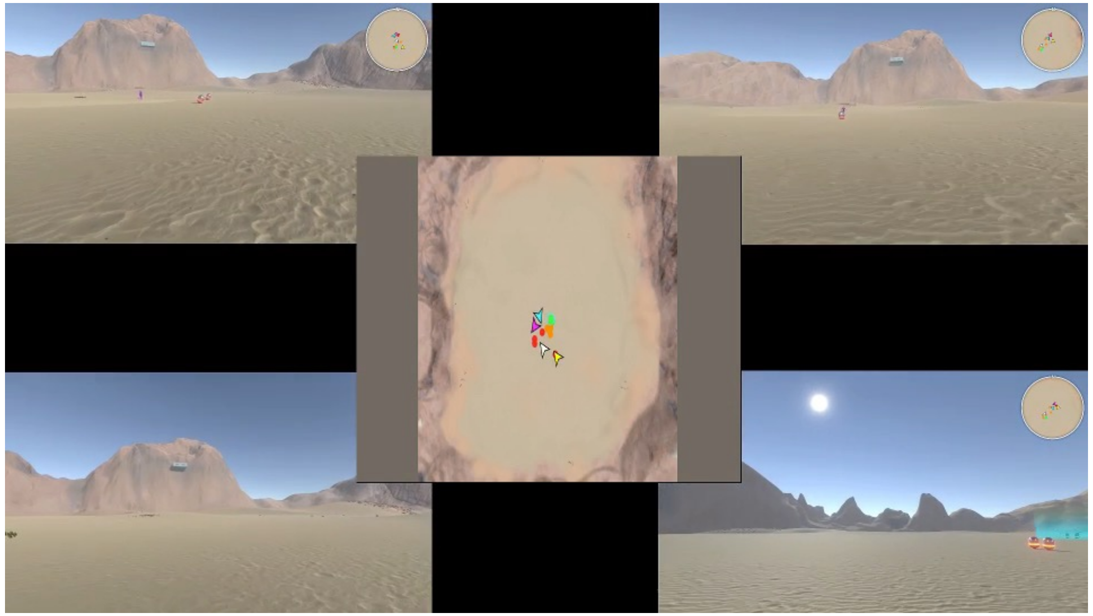

A. Bin Kamruddin, M J Richardson

- What about herding in more realistic 3D scenarios where navigation and exploration become essential parts of the problem?
- Or where the targets have more realistic dynamics (schooling fish, crowds etc)
- What about multi-agent reinforcement learning? How scalable?

University of BRISTOL

SSM

Thank you for your attention.

https://sites.google.com/site/dibernardogroup

MACQUARIE University SYDNEY-AUSTRALIA

MENTOR

F. Auletta

A. Lama

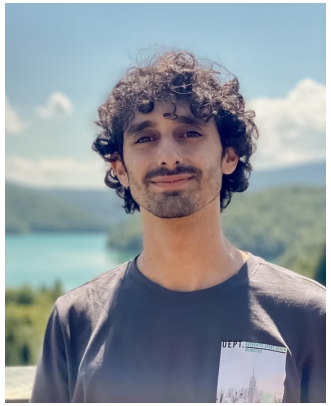

I. Napolitano

If you want to know more

- A. Lama, M. di Bernardo, “Shepherding Control and Herdability of Multiagent Complex Systems”, *Physical Review Research*, vol. 6, L032012, 2024
- F. Auletta, D. Fiore, M. J. Richardson, M. di Bernardo, "Herding stochastic autonomous agents via local control rules and online global target selection strategies", *Autonomous Robots*, 46(3), 469-481, 2022
- F. Auletta, R. W. Kallen, M. di Bernardo, M. J. Richardson, "Predicting and Understanding Human Action Decisions during Skillful Joint-Action via Machine Learning and Explainable-AI", *Scientific Reports*, 13, 4992, 2023
- F. Auletta, M. di Bernardo, M. J. Richardson, "Human-inspired strategies to solve complex joint tasks in multi agent systems", *Proc. 6th IFAC Conference on Analysis and Control of Chaotic Systems*, Catania, 2021