Shepherding with Multiple Sheepdogs Under Abnormalities

Edita Džubur and Adrien Doan

Collective Behaviour Course Research Seminar Report

November 16, 2025

Iztok Lebar Bajec | izredni professor | mentor

Introduction

Collective guidance, where one or more shepherding agents direct a flock toward a goal, has attracted interest across biological modelling, artificial life, and swarm robotics. Although herding appears simple in nature, reproducing it in artificial systems is challenging, as flock dynamics are nonlinear, interactions are local, and guidance must remain robust under uncertainty and potential agent failures.

Early shepherding research relied on biologically inspired heuristics, where complex group behaviour emerges from simple rules. These models successfully generate realistic flocking and herding patterns but may break down under noise or difficult initial conditions [2]. Control-theoretic approaches later reframed flock guidance as a containment or formation-shaping problem, offering stronger analytical foundations but often with reduced scalability or biological realism [3].

More recent work has shifted toward cooperative multi-shepherd systems. Multiple shepherds can provide better spatial coverage, but they also introduce challenges such as interference and a lack of automatic role allocation. These issues motivate decentralised mechanisms that enable coordinated multi-dog behaviour to emerge from local interactions.

Repulsive-force–based approaches offer one such mechanism. Kubo et al. show that introducing mutual repulsion between sheepdogs leads to stable circular formations and significantly improves multi-dog guidance efficiency [4]. Building on this idea, Tashiro et al. propose a related robust multi-shepherd model that enables dogs to coordinate around the flock and maintain effective guidance even under abnormal behaviours [1].

Our project aims to replicate and extend these frameworks. In this report, we outline the behavioural components of the model, summarise the interactions governing sheep and sheepdogs, and present initial implementation results that form the foundation for later experiments on multi-dog coordination and abnormal-agent resilience.

Methods

The shepherding model comprises two types of agents. Sheep (A-sheep), which move according to flocking principles and avoid dogs, and sheepdogs (s-dogs), which guide the flock toward a goal by repelling the sheep in a structured manner. The simulation evolves in discrete time, with each agent updating its state based on interactions with nearby agents.

A-sheep dynamics. Each sheep updates its acceleration as a combination of four behavioural tendencies: repulsion from nearby sheep to prevent overcrowding, alignment with the average velocity of visible neighbours, attraction toward the flock's local centre, and avoidance of dogs. These correspond to standard flocking terms, and may be expressed as a weighted sum

$$u_i = K_1 a_i + K_2 b_i + K_3 c_i + K_4 d_i,$$

where a_i, b_i, c_i, d_i denote the repulsion, alignment, attraction, and dog-avoidance components, respectively. The four gains K_1, K_2, K_3, K_4 determine how strongly each behavioural component influences a sheep's motion. The parameter K_1 regulates short-range repulsion and prevents overcrowding; K_2 sets the influence of alignment with neighbouring velocities; K_3 governs the attraction toward the flock's local centre; and K_4 adjusts the strength of avoidance when a dog is nearby. Together, these gains shape the characteristic flocking pattern of the A-sheep without requiring any detailed tuning for our preliminary implementation.

Figure 1 illustrates these interactions in diagrammatic form. Sheep consider neighbours within a fixed sensing radius, and each behavioural component depends only on local information. The resulting motion produces coherent flocking patterns in the absence of dogs.

Robust multi-agent shepherding under abnormal behaviours

Understanding how multiple autonomous shepherding agents can cooperatively guide a flock despite disturbances, failures, or misbehaving agents is important for both biological modelling and swarm robotics. The Tashiro et al. model introduces MSR-based filtering to tolerate anomalies among sheepdogs. Our work replicates this model and aims to extend its robustness and coordination strategies.

Collective behaviour | Shepherding | Swarm robotics | MSR algorithm | Anomalies

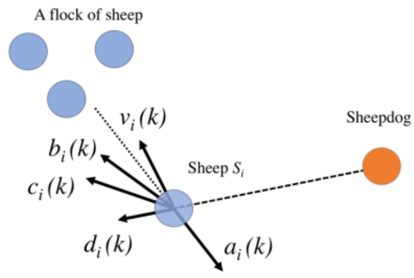

Figure 1. A-sheep behavioural interaction diagram.

Sheepdog dynamics. Each sheepdog acts on four behavioural forces. It is attracted toward the sheep farthest from the goal (the "target sheep"), experiences a short-range repulsion from that target to avoid collision, avoids getting too close to the goal itself, and repels other dogs to maintain spatial coverage. The dog's resulting velocity is given by

$$v_i = K_{f1}A_i + K_{f2}B_i + K_{f3}C_i + K_{f4}D_i,$$

where A_i encodes target pursuit, B_i short-range target repulsion, C_i goal avoidance, and D_i inter-dog repulsion.

The gains $K_{f1}, K_{f2}, K_{f3}, K_{f4}$ control the relative contribution of the four forces acting on each dog. The term K_{f1} determines how strongly a dog pursues the target sheep, while K_{f2} regulates the short-range repulsion that prevents collisions with that target. The parameter K_{f3} limits how close the dog can approach the goal region, and K_{f4} governs the repulsive interaction with other dogs, which promotes an even spatial distribution around the flock. These gains provide the behavioural balance between pursuit, spacing, and formation maintenance that characterises the s-dog model.

Figure 2 provides a conceptual illustration of the components acting on a shepherding agent.

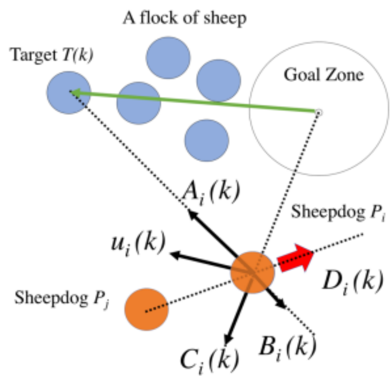

Figure 2. S-dog behavioural component diagram.

Abnormal behaviours and the MSR algorithm. Tashiro et al. [1] consider two forms of abnormality. In the first, actuation noise causes random fluctuations in the dog control gains at each timestep, modelling sensor or actuator faults. In the second, a subset of dogs is mis-programmed to steer the flock toward an incorrect goal. Both types degrade performance, and even a small number of abnormal agents can cause significant disruption.

To mitigate this, each normal dog employs a version of the MSR algorithm. At each timestep, a dog collects the states of neighbouring dogs, ranks them by similarity according to a chosen metric (spatial proximity, heading similarity, or distance to the goal), and discards up to f outliers before computing its cooperative action. This trimming eliminates the influence of abnormal agents, allowing the normal dogs to maintain coordinated behaviour.

Figure 3 summarizes the MSR filtering pipeline.

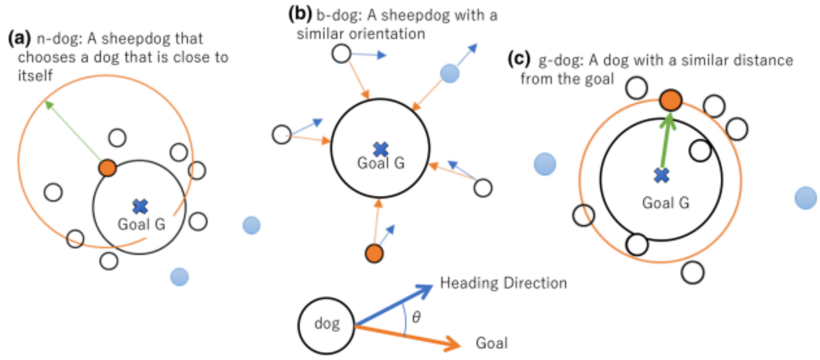

Figure 3. MSR algorithm flow diagram.

Simulation structure. Our implementation follows the discrete-time structure of Tashiro et al. At each simulation step, sheep update their positions using the flocking interactions described above, dogs select the target sheep and compute their four behavioural components, and MSR filtering is applied to the inter-dog interactions. Agents move with capped velocity, and interactions are limited by visibility radii. The simulation continues until all sheep reach the goal region or until a maximum time horizon is exceeded.

Results

Our first implementation milestone was to reproduce the baseline behaviour consisting of one dog and multiple sheep without abnormalities. Figure 4 shows an early screenshot from our simulation environment, demonstrating the expected target-selection behaviour and repulsive guidance of the flock. The system already exhibits coherent flocking and a well-defined pursuit of the most distant sheep.

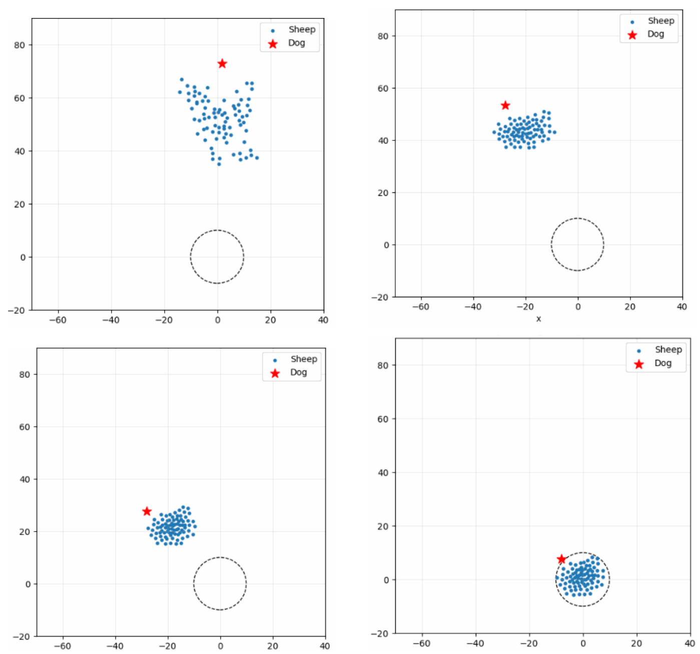

Figure 4. Preliminary implementation screenshots of our simple single-dog shepherding model. The dog (red star) drives the flock (blue agents) toward the circular goal region.

Although this first implementation omits MSR filtering and multi-dog interactions, it establishes the core structure required to extend the model. Position updates, flock dynamics, and goal-directed movement behave as expected. In the next stage of the project, we will incorporate multiple dogs and add the MSR filtering mechanism to assess robustness under abnormal conditions.

Discussion

The literature reviewed here reflects a progression from biologically inspired heuristics toward robust cooperative frameworks. Heuristic models can reproduce realistic herding patterns but lack resilience; control-theoretic approaches offer formal guarantees but do not scale well to decentralised or noisy environments. Multi-dog models provide improved coverage but introduce new vulnerabilities when some dogs fail or behave inconsistently.

The MSR-based framework addresses this challenge by filtering unreliable information, enabling normal dogs to ignore misbehaving neighbours. This property is particularly relevant for robotic systems, where sensor noise, actuator faults, and cyber-attacks are realistic concerns. Implementing and testing MSR in the shepherding setting therefore provides a valuable opportunity to study resilience in distributed multi-agent control.

Our next steps involve implementing multi-dog dynamics, integrating all three MSR variants, and experimentally evaluating robustness. We will compare performance across normal and abnormal conditions, quantify guidance time and success rate, and analyse how resilience scales with flock size, number of dogs, and the number of faulty agents. These experiments will constitute the next report.

1. M. Tashiro, M. Kubo, H. Sato, A. Yamaguchi, Guidance by multiple sheepdogs including abnormalities. *Artificial Life and Robotics* 27, 714–725 (2022).
2. D. Strömbom et al., Solving the shepherding problem: heuristics for herding autonomous agents. *Journal of the Royal Society Interface* 11, 20140719 (2014).
3. A. Pierson, M. Schwager, Controlling non-cooperative herds with robotic herders. *ICRA* (2017).
4. M. Kubo, M. Tashiro, H. Sato, A. Yamaguchi, Herd guidance by multiple sheepdog agents with repulsive force. *Artificial Life and Robotics* 27, 416–427 (2022).