Shepherding with Multiple Sheepdogs Under Abnormalities

Edita Džubur and Adrien Doan

Collective Behaviour Course Research Seminar Report

January 11, 2026

Iztok Lebar Bajec | izredni professor | mentor

This report presents the third stage of our Collective Behaviour project, focused on extending a replicated multi-sheepdog shepherding model with abnormal agent behaviours and environmental obstacles. Building on the robust guidance framework introduced by Tashiro et al. (2022), we implement sheepdogs with abnormal inputs and objectives, incorporate static obstacles into the environment.

Introduction

Collective guidance, where one or more shepherding agents direct a flock toward a goal, has attracted sustained interest across biological modelling, artificial life, and swarm robotics. Although herding appears simple in natural systems, reproducing it in artificial environments remains challenging due to nonlinear flock dynamics, local interactions, and sensitivity to disturbances or agent failures. These challenges become more pronounced in multi-shepherd settings, where coordination among guiding agents is required to achieve stable and efficient guidance.

Previous work has shown that cooperative multi-dog shepherding can significantly improve performance when appropriate interaction mechanisms are employed. In particular, introducing repulsive forces between sheepdogs enables self-organised formations around a flock, reducing oscillatory behaviour and overshoot. Building on this idea, Tashiro et al. proposed a robust multi-shepherd framework that integrates Mean Subsequence Reduced filtering to tolerate abnormal sheepdog behaviours such as actuation faults or incorrect objectives.

In earlier stages of this project, we replicated the A-sheep and s-dog models and validated multi-dog shepherding under nominal conditions. In the present stage, we extend the system in two directions. First, we explicitly implement abnormal sheepdog behaviours and examine their effect on multi-dog coordination using MSR-based filtering. Second, we introduce static obstacles into the environment in order to observe shepherding behaviour in constrained spaces. The emphasis of this report is therefore placed on behavioural outcomes and experimental observations rather than on re-deriving the underlying models.

Methods

The simulation consists of two agent types, sheep and sheepdogs, evolving in discrete time within a bounded two-dimensional workspace. Sheep follow flocking dynamics with local interactions, while sheepdogs influence the flock indirectly through repulsive guidance toward a goal region. The base A-sheep and s-dog models follow the formulation of Tashiro et al. and were implemented and validated in earlier stages of this project. In this section, we provide a concise overview of these dynamics and describe the extensions introduced in the present stage.

Sheep dynamics. Sheep agents follow flocking dynamics based on local interactions with nearby agents and sheepdogs. Their motion is governed by a combination of short-range repulsion from neighbouring sheep, which prevents overcrowding, alignment with the average velocity of visible neighbours, which promotes coherent movement, and attraction toward the local centre of the flock, which maintains group cohesion. In addition, sheep actively avoid nearby sheepdogs, introducing directed motion away from regions of high dog influence.

All interactions are local and depend only on agents within a fixed sensing radius. The relative influence of repulsion, alignment, attraction, and dog avoidance determines the emergent flocking behaviour and shapes how the flock responds to external guidance. These mechanisms together produce coherent flock motion while allowing rapid reconfiguration in response to sheepdog positioning.

Sheepdog dynamics. Sheepdogs guide the flock indirectly by positioning themselves relative to the sheep and the goal region. At each timestep, a sheepdog identifies a target sheep, defined as the individual farthest from the goal, and moves so as to exert repulsive pressure that encourages flock motion toward the desired direction. To avoid

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

collisions, sheepdogs maintain a short-range separation from their target and from other dogs.

In addition to target-oriented motion, sheepdogs avoid approaching the goal region too closely and repel one another to maintain spatial coverage around the flock. This inter-dog repulsion plays a key role in preventing multiple dogs from converging on the same region and promotes the emergence of a surrounding formation during guidance. The balance between target pursuit, collision avoidance, and inter-dog spacing determines how effectively sheepdogs coordinate their actions.

Abnormal sheepdog behaviours and MSR filtering. Two types of abnormal sheepdog behaviour are implemented. In the first case, abnormal inputs are modelled by introducing stochastic perturbations to sheepdog control gains at each timestep, representing sensor or actuator faults. In the second case, abnormal objectives are introduced by assigning a subset of sheepdogs an incorrect goal location, causing them to actively guide the flock in a conflicting direction. These abnormal behaviours persist throughout the simulation and directly affect local coordination among sheepdogs.

To limit the influence of abnormal sheepdogs, normal agents apply Mean Subsequence Reduced (MSR) filtering when processing information from neighbouring dogs. At each timestep, each sheepdog ranks neighbouring agents according to a similarity metric derived from observed state variables and excludes a fixed number of the most dissimilar neighbours before computing its cooperative response. This local filtering mechanism allows normal sheepdogs to attenuate the impact of anomalous behaviour without requiring explicit identification of faulty agents.

The MSR-based filtering process and its role in inter-dog interaction are illustrated schematically in Figure 1.

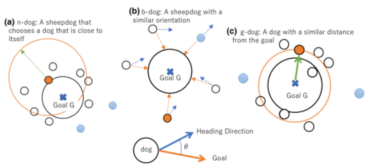

Figure 1. Schematic illustration of MSR-based neighbour filtering. Adapted from [1]

Obstacle incorporation. Static obstacles are incorporated into the simulation environment to constrain agent motion and introduce additional environmental complexity. Obstacles occupy fixed regions of the workspace and must be avoided by both sheep and sheepdogs during guidance. Agents account for obstacle presence during motion updates, ensuring collision-free trajectories while maintaining shepherding behaviour.

For sheepdogs, obstacle avoidance is implemented through a geometric correction of the intended control acceleration. When a sheepdog is predicted to move toward an obstacle $\vec{v} \cdot \vec{n} > 0$ (dog going towards the obstacle), its acceleration vector is decomposed into components normal and tangential to the obstacle boundary. If the normal component points into the obstacle, it is removed and only the tangential component is applied, allowing the sheepdog to slide along the obstacle surface rather than penetrating it. This projection-based correction avoids introducing additional repulsive force terms and preserves the structure of the original guidance dynamics in free space.

The obstacle-aware acceleration correction is illustrated schematically in Figure 2.

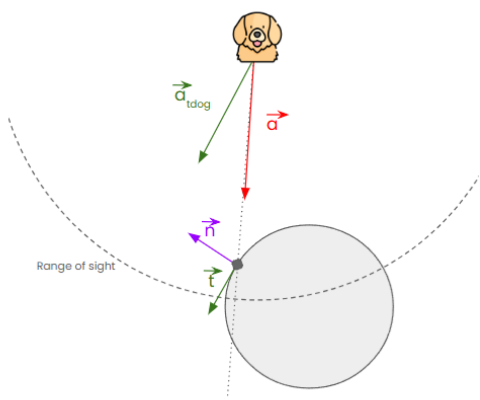

Figure 2. Obstacle-aware acceleration correction for sheepdogs.

Simulation setup. Simulations are initialised with sheep randomly distributed within a bounded region and sheepdogs positioned outside the flock. Each run proceeds until all sheep reach the goal region or a maximum time horizon is exceeded. Multiple runs are performed with different random initial conditions to account for stochastic effects introduced by abnormal inputs and agent placement. Guidance performance is evaluated using success rate and average guidance time across runs, together with qualitative observations of formation stability and flock cohesion.

Results

Shepherding performance was evaluated under four scenarios designed to isolate the effects of multi-agent coordination, abnormal sheepdog behaviour, and environmental constraints. Simulations were run until the flock reached the goal region or a maximum time horizon was exceeded.

Results are presented as representative temporal snapshots illustrating typical formation structure, flock cohesion, and guidance dynamics observed across runs.

Single sheepdog. As a baseline, we first consider shepherding with a single sheepdog. Guidance is achieved by positioning behind the flock relative to the goal and applying repulsive pressure to induce motion. While successful in most runs, guidance is sensitive to flock geometry and often exhibits oscillatory behaviour as the dog repeatedly repositions to correct stragglers. The absence of cooperative support limits the ability to stabilise flock shape, particularly as the number of sheep increases.

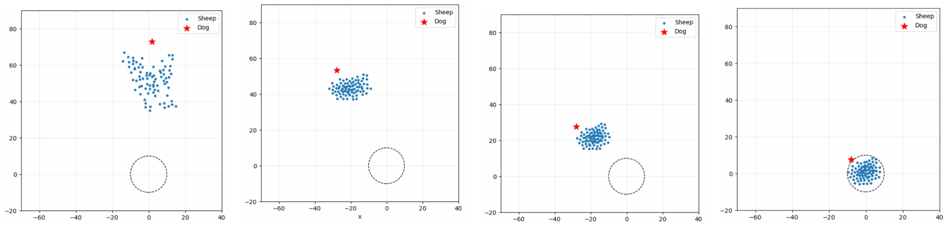

Figure 3. Temporal evolution of shepherding with a single sheepdog.

Multiple sheepdogs. Introducing multiple sheepdogs significantly improves guidance performance. Inter-dog repulsion leads to the emergence of a surrounding formation, allowing the flock to be guided more smoothly and with reduced oscillation. The distributed pressure applied by multiple agents helps maintain flock cohesion and accelerates convergence toward the goal. Effective multi-dog guidance requires careful parameter tuning to avoid oscillatory motion and unstable formations.

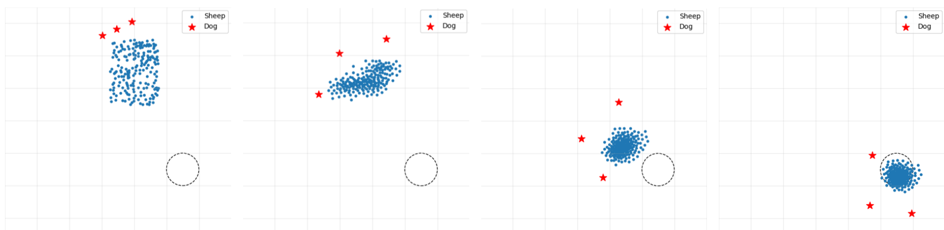

Figure 4. Temporal evolution of nominal multi-dog shepherding.

Multiple sheepdogs with abnormalities. In this scenario, a subset of sheepdogs exhibits abnormal behaviour, either through perturbed control gains or incorrect goal objectives. These agents introduce local disturbances and occasionally disrupt the surrounding formation. With MSR-based filtering enabled, normal sheepdogs are generally able to reduce the influence of abnormal agents and re-establish effective guidance, although convergence toward the goal is slower and formation symmetry is temporarily degraded. Abnormal sheepdogs intermittently occupy structurally important positions around the flock, requiring normal agents to continuously re-adapt their formation through local correction.

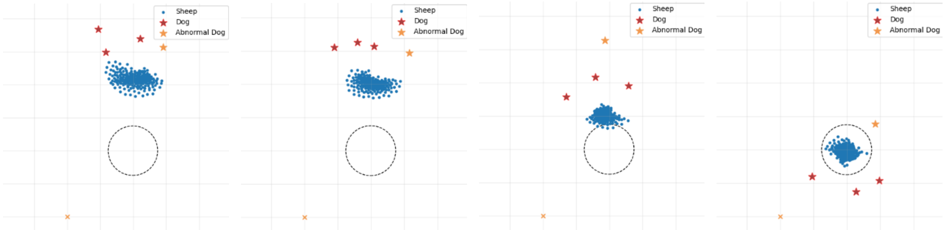

Figure 5. Temporal evolution of multi-dog shepherding with abnormal sheepdogs.

Multiple sheepdogs with obstacles. Static obstacles introduce environmental constraints that affect both flock motion and sheepdog coordination. As sheepdogs manoeuvre around obstacles, the surrounding formation becomes temporarily distorted, particularly in narrow regions. Despite these constraints, guidance remains successful in most runs, with sheepdogs adapting their positions once sufficient free space becomes available. Obstacles compress available steering space, temporarily reducing formation symmetry and limiting the effectiveness of inter-dog repulsion in maintaining coverage around the flock.

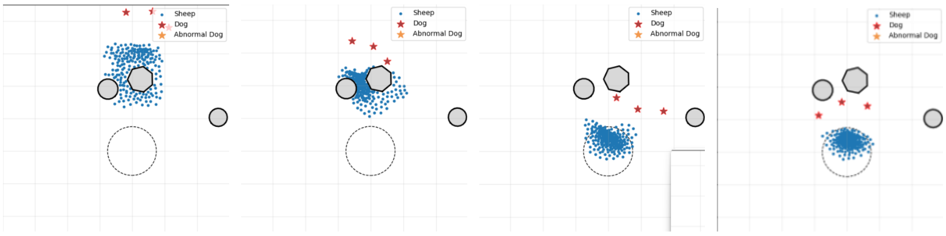

Figure 6. Temporal evolution of multi-dog shepherding in the presence of static obstacles.

Discussion

The results demonstrate that cooperative multi-dog shepherding remains feasible, although performance is strongly sensitive to parameter choices. Abnormal sheepdogs primarily affect guidance through unstable motion or persistent misdirection, while obstacles disrupt formation geometry and reduce the available steering space, increasing the difficulty of maintaining coordinated motion. In unconstrained environments, the observed behaviour is broadly consistent with the reference model of Tashiro et al., where inter-dog repulsion supports surrounding formations and MSR-based filtering limits the immediate influence of anomalous agents. However, effective guidance in these settings requires careful tuning of control gains and filtering parameters, and robustness emerges only within a relatively narrow parameter regime, with normal sheepdogs relying on repeated local correction to compensate for disturbances.

The introduction of obstacles reveals behaviours not addressed in the original work and further exposes the sensitivity of the system to parameter selection. Spatial constraints compress formations, reduce symmetry, and create bottlenecks where sheepdogs compete for limited steering space, leading to temporary distortions of the surrounding configuration and increased guidance time, even when all sheepdogs behave nominally. These findings highlight a limitation of purely local coordination mechanisms in constrained environments. While MSR-based filtering can attenuate anomalous behaviour in open settings, it does not address challenges introduced by spatial constraints, suggesting that obstacle-aware coordination, adaptive parameter adjustment, or higher-level formation control may be required for reliable guidance.

CONTRIBUTIONS. E.D. implemented the single-dog shepherding simulation and established the overall software architecture. A.D. extended the implementation to the multi-dog setting, integrated abnormal behaviours, and incorporated obstacle handling into the simulation environment. For the written work, E.D. prepared and refined the reports for subsequent stages, ran simulations for the report visualisations, and performed parameter fine-tuning.

References

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).