How the engineering of Swarming Behaviors could be drawing on Social Sciences

Hanno Hildmann[1,2,*], Regina Nockerts[3], Sebastian Coffeng[4], Hannu Karvonen[4], Johan van der Heuvel[1], Ricardo van der Pluijm[5], Fabrice Saffre[4]

[1]Netherlands Organisation for Applied Scientific Research (TNO), Autonomous Systems & Robotics Group, Oude Waalsdorperweg 63, 2597 AK Den Haag, The Netherlands - {hanno.hildmann; johan.vandenheuvel}tno.nl
[2]The Hague University of Applied Sciences (THUAS), Smart Sensor Systems Research Group, Department of Mechatronics, Rotterdamseweg 137, 2628 AL Delft, The Netherlands
[3]University of Amsterdam (UvA), School of Computational Social Sciences, Roeterseiland Campus, Roetersstraat 11, 1018 WB Amsterdam, The Netherlands - r.a.nockerts@uva.nl
[4]Technical Research Centre of Finland Ltd. (VTT), MIKES bldg, Tekniikantie 1, 02150 Espoo, Finland - {fabrice.saffre; hannu.karvonen; sebastian.coffeng}@vtt.fi
[5] Royal Netherlands Aerospace Centre (NLR), Department Aerospace Training and Simulation, Anthony Fokkerweg 2, 1059 CM Amsterdam, The Netherlands - ricardo.van.der.pluijm@nlr.nl

Keywords: Emergence, behaviors, complex systems, system analysis, robotic systems, swarms.

Abstract

Emergent behavior in swarming and self-organizing systems can be systematically studied. This paper proposes to draw on excisting conceptual frameworks from the social sciences to do so. To this end, the article proposes a theoretical framework that treats local agent behaviors as "policies" (a concept borrowed from social sciences) and uses insights from policy analysis and systems analysis (established practices in social sciences) – such as feedback loops, institutional constraints, and emergent norms – to inform the design of robotic swarms (specifically: homogeneous swarms of autonomous platforms operating in 2D with communication constraints). To illustrate this, two canonical tasks, *Area Exploration* and *Area Coverage*, are used as case studies. By examining these tasks through a social science lens, we illustrate how mechanism design principles, phase transitions, and decentralized coordination strategies contribute to robust emergent dynamics. This has its practical limitations and we discuss these as well.

1. Introduction

The study of emergent behavior – how large-scale patterns arise from local interactions – has been studied across disciplines ranging from ecology and sociology to robotics (Holland, 1995). We foresee that cooperative collectives composed of drones or UGVs will soon be deployed as swarms, where simple agents following basic rules (Beckers et al., 2000) can produce complex, adaptive group dynamics (Bonabeau et al., 1999). In this novel field, designing and modelling swarm behaviors is drawing on biology (e.g., insect colonies; Dussutour et al. 2004) and using heuristic approaches. We propose that conceptual frameworks from the social sciences, especially policy analysis and systems analysis, offer a rich and underutilized source of insight for understanding and engineering swarm dynamics.

We know that simple local protocols (e.g. neighboring agents following basic interaction rules (Schelling, 1971)), can enable individual agents to cooperate based solely on local information without any centralized control (Beckers et al., 2000). These interactions are continually shaped by feedback signals (such as reward or reinforcement cues that inform agents about the success of their actions; Sutton and Barto 1998), allowing each agent to adjust and refine their behavior over time. Through repeated local interactions and ongoing adaptation (driven by feedback), coherent emergent conventions eventually form at the group level such as e.g., the emergence of consistent social norms or coordination schemes solely from repeated (individual and decentralized) interactions (Young, 1993).

This combination of local rule-following and feedback-driven adaptation enables multi-agent systems to achieve organized, global outcomes; and to do so from the bottom up.

* Corresponding author

The idea that macro (system-level) outcomes depend on micro (device-level) interactions is found in the work by e.g., (Holland, 1995) and (Miller and Page, 2007). Complex global behaviors emerge seemingly spontaneously from the interactions of relatively simple, adaptive agents. They highlight the concept of emergence (Johnson, 2002) and (Minati and Pessa, 2006), where macroscopic outcomes arise without central control and in the absence of explicit global goals. Key theoretical insights include mechanisms such as feedback loops, adaptation, and self-organization: agents continually adjust their behavior based on local information, often guided by simple rules or heuristics. Through this the entire system can converge toward robust and sometimes surprising global patterns, which Holland referred to as some *hidden order* (Holland, 1995). Basically, individual interactions lead to system-level phenomena through dynamic mechanisms such as tipping points and equilibria.

Ostrom's institutional governance framework, termed *governing the commons* (Ostrom, 1990), investigates how groups of individuals can self-organize to manage shared resources in the absence of a central authority: micro-level behaviors, guided by locally defined institutional rules (such as clearly defined resource boundaries, monitoring compliance, and graduated sanctions), can yield robust macro-level governance structures. Her theories stress collective action, social norms, and institutional arrangements, which act as explicit or implicit contracts guiding the behaviors of individuals. Through such rules one can foster cooperation, prevent free-riding, and encourage sustainable resource use. She showed how institutions shape incentives and, through these, drive emergent social outcomes. Ostrom's framework is rooted in empirical studies of real-world collective action problems, it highlights the importance of rule compliance, local adaptability, and community-level decision-making.

The key difference lies in the emphasis: Holland and Miller focus more broadly on decentralized adaptive interactions and emergent complexity without necessarily embedding explicit rules or formal institutions while Ostrom explicitly addresses how particular institutional arrangements, rules deliberately devised by participants, facilitate desired collective outcomes.

The work of (Holland, 1995) and (Miller and Page, 2007) considers how complex propeties emerge *naturally* from simple interactions (often based on natural laws) while (Ostrom, 1990) looks at how human-designed institutional rules and social norms can give rise to stable social group behaviours.

Common to all of them is the insight that macro-level outcomes depend fundamentally on micro-level interactions. Whether these system level behaviours are framed as system robustness, resilience, or effective governance simply depends on the lense through which we look at the system. Both fields have contributed to the understanding of emergence and collective behaviours and we argue that, so far, Computer Scientists working with robotic collectives (swarms) have largely ignored the works by prominent practitioners in the Social Sciences.

1.1 Scope, overview and contributions

In this paper, we develop a theoretical framework grounded in these cross-disciplinary connections. We treat each agent's controller or behavior rule as a "policy" in the social sense (§2.1), subject to evaluation and possible adaptation. We then use concepts such as feedback loops (§2.2), institutional constraints (§2.3), emergent norms (§2.4), mechanism design (§2.5), phase transitions (§2.6) and robustness (§2.7) to link the agents' policies to the collective outcome.

Two concrete swarm missions illustrate the approach: **(1) area exploration** (minimizing time to detect a target, cf. §3.1) and **(2) area coverage** (minimizing the average time since each region was last visited, cf. §3.2). These common tasks nicely illustrate the contrast between either focusing on a one-time collective achievement or engaging in an ongoing collaborative effort. We will use this to highlight how different policies and coordination mechanisms are needed.

Our analysis draws on established ideas from complexity science and robotics while emphasizing original connections to social theory. We discuss how intentional mechanism design at the agent level can induce robust self-organization at the system level, including managing phase transitions in swarm behavior and ensuring resilience to disturbances.

2. Conceptual Framework: Policies and Emergent Mechanisms

We start by discussing the idea underlying the framework, which is to see the behaviors of agents as determined by policies similar to how the actions of an individual in society are governed by strategies or rules. This already makes the connection between social sciences and swarm engineering that we see clear. When looking at this from a system's perspective a swarm (of either robots of humans in a society) is a complex adaptive system as described by (Holland, 1995) (Miller and Page, 2007). In this systems, the interactions between the individual agents, paired with the individual agent's ability to act in response to the perceived actions of others, can give rise to emergent phenomena (for example, the efficient exploration of an area).

Our goal is to asses whether approaches to complex system control from social sciences can be applied to swarm engineering.

2.1 Agents as policy implementers

In social sciences there is the concept of a *policy*, referring either to a set of rules or some benefits or incentives designed to influence the behavior of individuals. This is akin to the control mechanism (often an algorithm) in robotic swarms that control the actions of the individual agent. For example, the rule *"move to the nearest region that has not been visited recently"* could be such a policy. Policy analysis is a practice in social sciences which can be translated to the engineering of swarms as a tool to predict the outcomes of such rules and the changes therein. General policies can be tuned to impact the agents by e.g., changing the parameter that determines the extend to which an agent interprets the word *recent* in the rule stated above.

By formulating agent behaviors as policies, we enable the use of policy analysis tools to predict and evaluate swarm outcomes. Each agent's policy can be tuned as if it were a lever in an institutional design problem: we can ask how changes to the policy parameters (e.g., how strongly an agent avoids recently visited areas) affect collective performance measures.

2.2 Feedback loops and system dynamics

A fundamental aspect of both, swarms as well as social systems is the feedback loop. The agents'ability to react to changes in their environment (and thus to the actions of their peers) is integral to the concept of self-organization and emergence. Feedback can come as positive (in which case it serves to reinforce some rule or policy) or negative (to balance out the system): in a swarm, an example for a positive feedback is an event that triggers a rule to attract others to a location while seeking distance from others due to overcrowding is an example of negative feedback. examples can be applied to each: natural swarms such as ants, human societies as well as robotic collectives.

We can include feedback mechanisms by design: by ensuring that a specific event triggers a broadcast to attract nearby peers (positive, e.g., (Saffre et al., 2022a)) or by emitting a stress signal that correlates with the perceived overcrowding around an agent and which repulses nearby agents (negative, e.g., (Saffre et al., 2023)). When both, positive and negative feedback are used in interconnected mechanisms (Hildmann et al., 2019) then the resulting swarm behaviors can be designed to alternate or occur only when the situation and the environment require it.

2.3 Institutional constraints and rules of engagement

In social sciences, institutions are seen as a set of shared rules or constraints to govern interactions. According to (Ostrom, 1990) such institutions are a driving factor in societal self-organization. This also translates directly to swarm engineering: we may want to impose such institutional constraints to guide and facilitate operational aspects such as the need to navigate in the same environment (and therefore have a need for some sort of obstacle avoidance protocol). A straight forward example is the side of the road on which we drive. Ultimately it does not matter whether this is the left or the right, as long as everyone adopts the same as their default. In a society well-designed (or evolved) rules can improve overall efficiency and prevent collective failures and the hope is that these results can be transferred to swarm engineering.

The impact that can be achieved through the monitoring of compliance or by finding a way to sanction the behavior of agents that deviate from the norm (Ostrom, 1990) can form the basis of a design principle, where it is used to enforce a specific group behavior or property. Analogous mechanisms in swarms could be used to self-regulate e.g., the consumption of resources.

2.4 Emergent norms and conventions

In addition to clearly stated formal rules there are often also *unwritten* rules in societies. These rules do not exist from the beginning and are thought to emerge over time and due to the existing formal rules. These emergent norms arise through repeated interactions in a group (Axelrod, 1986). In social groups this also helps to overcome the individual quirks a single specific agent may exhibit. While such individual traits are largely absent from robotic swarms (where systems of the same class and type can often be considered to be identical) they can still develop such conventions. An easy example for this is the self-organizing of a swarm to provide coverage over an area. This can be achieved by enabling the platforms to partition the area among themselves, effectively creating *territories* but doing so without central coordination. A norm here could be to avoid trespassing on some other agent's territory, which at system level improves efficiency of the swarm by reducing redundancy.

There are models for this emergence of norms (Axelrod, 1986) in the literature and we can use them to design system where specific conventions are *encouraged*. This can be achieved through some degree of learning, for example when agents have some memory and can remember the outcome of previous actions or interactions. When including this memory in the decision making process a well designed system can be gently guided towards adopting rules that *seem to work*, resulting in a swarm wide norm that no individual agent planned (and indeed, no individual agent may be aware of as such).

2.5 Linking micro to macro: the mechanism design lense

In the field of economics, the theory of mechanism design focuses on the design of rules (i.e., the mechanisms) that aim to achieve system level outcomes by governing the individual agents' actions. The underlying assumption is that the agents are acting rational, meaning that their decision process can be modeled using some notion of rationality. From these, the desired collective outcomes can be constructed (Hurwicz and Reiter, 2006). In human society the rational is often that one will chose an action that leads to preferable outcomes, where what is preferred is rather subjective (but can be modeled at societal level). Robots do not have subjective preferences as such but they do have clearly programmed policies and rules and they do normally have an objective function to assess the state of the environment as well as their own. Nevertheless, we can design the rules governing the local interaction with others (the mechanism) to converge the swarm behavior (the emergent outcome) towards some desired property.

For example, we can drive the resource / task allocation within a swarm by including some capacity based weighting in the decision process. In (Hildmann et al., 2012) the re-allocation of clients to servers is stochastic but the probabilities are weighted to reflect the capacity of the agents participating in the exchange / the competition. An important property of this is the extent to which the resulting solution is stable and at equilibrium.

By designing an appropriate (local) notion of utility or preference each agent can be steered to (locally) optimize their own performance but through that achieve some (global) near-optimum and optimize system performance. This is similar to how (Hurwicz and Reiter, 2006) designs incentives in order to ensure that the (game-theoretic) Nash equilibria falls in what the social optimum / socially preferred optimum is.

2.6 Phase transitions and critical points

A phase transition happens when a small change in the parameters of a system *pushes* the system past some (imaginary) threshold, resulting in a qualitative change in behavior. The seemingly coordinated attack on a single target or location by an army of ants or in packs of mammals the seemingly sudden decision to flee are good examples. (Schelling, 1978)'s example is the sudden onset of cooperation in a social setting.

Robotic swarms also undergo transitions: reducing the communication range of drones in a swarm will likely lead to monotonous performance reduction until at one point the swarm *falls apart* and seizes to operate as a whole. From this example its easy to see that understanding such critical points and the parameters that lead to phase transitions can be extremely important. We therefore suggest and recommend analyzing the behaviors of swarms as functions of specific parameters (such as e.g., agent density, the aforementioned communication range, noise levels) to identify the *tipping points*. Around these points (in parameter space) small changes in parameter values can lead to fundamentally different system outcomes and behaviors.

In Schelling's model (Schelling, 1978) a slight increase in intolerance results in complete segregation. Armed with the knowledge of this as well as the insight into the parameter values that trigger this, we can then set out to either design policies that avoid this (parameter space) altogether or leverage them for emergency behavior in extreme circumstances. We can use insights from statistical physics or from network sciences from social modeling to design a mechanism that e.g., intentionally triggers the swarm to engage in defensive operations.

2.7 Robustness and self-repair

Swarms and collectives, be they human, animal or robotic, can be expected to *live* in a noisy world, full of uncertainties, impurities and disturbances. Societies mitigate the impact of such disturbances by sticking to broadly stated rules and through individuals adapting their behaviors. This provides resilience and allows societies to persist through tumultuous times.

We would like to make the robotic swarms robust and resilient as well. This means to design the rules governing the agents to be redundant (this is to allow multiple mechanisms to guide specific aspects, so that if one fails another can kick in). For example, agents can be programmed to identify and detect failure in their peers (by maintaining a memory of when one last saw a peer or by corroborating information through other channels) and then adjust their policies to bring about, at system level, a behavior that mitigates the issue. Examples are a swarm that extends its coverage (by decreasing local density / by increasing the preferred distance between individual agents) to re-connect to a lost swarm member. In a society this happens when someone briefly takes up the tasks of someone else to *"keep the show running"*. This can happen decentralized and without central triggers. The ability to self-repair is akin to the robustness of social-ecological collectives (Ostrom, 1990).

We propose to test swarm policies under various settings and scenarios where the swarm will fail. This is akin to the policy stress-testing when designing laws and policies. Only robust behaviors will continue to emerge (at the system level) despite some individual agents failing (at agent level). We therefore combine the agent-level policy design with system-level mechanism analysis. We use feedback loops, institutional constraints, and norms to design systems and anticipate on phase transitions to engineer swarm behaviors in robotic systems.

3. Emergent design for exploration vs. coverage

A swarm's required emergent behavior can differ significantly between missions. We compare two commonly studied scenarios: (1) one-time area exploration, cf. in Section 3.1, where the goal is to detect targets or map an unknown environment as quickly as possible, and (2) persistent area coverage, cf. Section 3.2, where the goal is to continually visit every region so that no location goes unobserved for long. While these tasks sound similar (both involve spatial distribution of agents), they pose different demands on coordination and thus benefit from different agent-level policies. Using our framework, we treat each scenario as a case study in designing local policies (and the implied social mechanisms) for the desired global outcome.

3.1 Area exploration: minimizing detection time

In area exploration, a team of robots must collectively explore an unknown or uncertain space to find a target. The performance is often measured by the *detection time* – how quickly at least one agent finds the target. This is analogous to a group of individuals searching for a critical item or survivor in a disaster area. The emergent objective is a rapid, exhaustive search.

Therefore, we want to maximize the rate of exploration (rapid) and at the same time ensure that this extends to all locations (exhaustive). By doing so we decrease the time to detection. A simple, but sub-optimal, strategy is to let the agent's perform a random walk. Given enough time, this will eventually find the target but it will do so with a large amount of redundancy (locations visited more than once) (Burgard et al., 2005). Agents should coordinate their actions and generally avoid overlap in their search areas. One simple rule to achieve this would be : *"if you sense another agent nearby, move away to spread out"*. This rule is found in animal groups (Couzin et al., 2005) where it causes a group to distribute somewhat evenly over an area. We already mentioned something like this as an example of a negative feedback loop in Section 2.2.

Another insight is that restricting communication is likely to have a deteriorating effect on the swarm performance. In most cases, the ability to share information is beneficial but realistically we cannot expect communication to be infinite and pervasive (Husain et al., 2022). It is therefore advisable to design policies that can continue to function under contested communication and that operate mainly (or entirely) on locally available information. The *frontier exploration* heuristic (Burgard et al., 2005) is an example for this: agents move towards the nearest boundary (where visited locations meet unvisited locations), thereby creating a swarm that expands outwards to rapidly explore a region of space. If we enable agents to share their maps with other agents in their proximity this process can be achieved entirely decentralized: the resulting emergent behavior is coordinated expansion without central control.

In addition to rapidly exploring an area we might also want to exploit our performance, i.e., to do something once we found the target. This is a good example of a positive feedback, because agents should communicate when a target is found (and where) and thereby recruit other agents to this location. The benefit in doing so is the shortening of the time until the other agents become aware of the target (in other words, we save the time they would otherwise spend looking for a target in all the wrong places).

A variation of this scenario is where the target is already detected but not yet confirmed as to do so requires multiple agents to cooperate (let's say because their individual sensor arrays are by themselves not sufficient to penetrate the target's electronic shielding). In this variation there is a temporary need to attract local peers (akin to the rapid response in social system to trhe shout *"help"*) but once the issue is resolved the other agents can return to their previous tasks. This temporary phase transition would be driven by a positive feedback (the detection of the target by a nearby, and thus within communication range, agent) but this would be short lived as it would soon be offset by a negative feedback when the issue is resolved. Once there is no identification task available the density of the agents exceeds their preferred value and they disperse again, resulting in a second phase transition.

With care we can also provide robustness to this to avoid over reaction or to come running for a false alarm. For example, more than one notice of detection could be required to converge on the location, or agents could weigh their willingness to be recruited by e.g., the number of already recruited other agents. The cost of this would be the response time but this is a trade-off that can be evaluated and designed. This also has parallels with social sciences: emergent group attention (*"help"*) and signal filtering (two alarms needed before the agent reacts) are found om social information sharing networks.

The idea to apply this to robotic swarm research is not entirely new of course: Burgard et al. (Burgard et al., 2005) demonstrated coordinated multi-robot exploration through letting the agents pick their targets and search regions, i.e., some sort of implicit negotiation for tasks. When looking at this through the policy lense we include in the agent's objective function that it prefers to explore currently unknown locations and that it is penalized for revisiting locations that are already known. Enabling individual agents to optimize to their preferences, the entire swarm improves their detect and identify performance.

For exploration tasks the recommended agent policies are:

1. **Dispersion and division of labor:** avoid duplicating other agents' coverage areas (an emergent norm of dispersing).
2. **Information sharing:** whenever possible, communicate discoveries or map updates to peers (an institutional rule enabling collective knowledge).
3. **Responsive clustering on detection:** use feedback to adapt behavior when a likely target is found – agents converge on locations (a positive feedback akin to an alarm response), but use checks to avoid misinformation.

Under these policies, the swarm exhibits emergent behavior that closely matches the desired outcome: broad initial coverage with rapid focus on true targets. The performance approaches that of an optimal, centralized, system (Burgard et al., 2005) but is really achieved through decentralized, robust means.

https://doi.org/10.5194/isprs-archives-XLVIII-2-W11-2025-119-2025 | © Author(s) 2025. CC BY 4.0 License.

3.2 Area coverage: minimizing revisit time

Area coverage involves patrolling or monitoring a region indefinitely, such that the time since any point was last visited by an agent is minimized (or bounded). This is crucial for e.g., information surveillance and reconnaissance (ISR), environmental monitoring, and security patrols. Unlike exploration, coverage has no termination criteria that ends the process; it is about persistent attention to an area, and that as uniformly as possible. The emergent objective is a sustained pattern that optimally distributes visits over space and time.

Whatever the strategy chosen, it must ensure that a member of the swarm visits each location with a reasonable frequency to ensure a high performance quality (i.e., minimizing the time since the the last visit to any location) while avoiding visiting locations recently visited by a peer (i.e., avoiding redundancy and inefficiencies). A simple version of such a strategy somehow partitions the area in individual territories, and assigns each on as the responsibility for a specific swarm member. This approach, which is found in e.g., bees partitioning the area around the hives in sectors, makes it easy to avoid redundancies. A simplistic rule that could be designed to govern this territorial behavior could be *"if you encounter another agent's recent path or marker, turn away"*, which is similar to using the stigmergy mechanism where e.g., pheromones are used to indicate a path was recently traveled can be used to influence the movement of other swarm members (Bonabeau et al., 1999). Combined with some memory of where one was and where one would like to travel again, this can already lead to some partitioning of the area and result in individual agents *patrolling* their territory, following their own traces and avoiding those of their peers. In a way, this could be visualized as *tagging* the environment to continuously claim one's beat.

Among the approaches used to partition areas are Voronoi tessellations (Okabe et al., 2000) which essentially compute optimal partitions (Cortés et al., 2004) of an area subject to properties or capabilities of the individual areas / resources allocated to the area. The way this is achieved is not by computing the optimal division lines but instead by constructing sub-optimal ones and then iteratively amending them based on some intricate interplay between only bordering regions. For example, if two bordering areas differ greatly in size then the one that is larger is likely to allow the shift of the border slightly inside its territory, thereby moving a small bit towards equalizing the areas in the two bordering territories. The outcome, as shown by (Cortés et al., 2004), is efficient coverage. Furthermore, by the nature of the approach, this outcome is dynamic and will be continuously updated and improved which means that if outside factors result in changes to the overall problem (e.g., fewer swarm members, a smaller or larger area, etc) the swarm will continue to improve the (now sub-optimal) solution until it is optimal again. In our example we could calculate a performance value for each territory (to represent the average time it has been since any location has been visited) and use those to amend the borders: the territory with the higher value (and thus worse performance) is reduced with parts being taken over by an agent with a better performing territory.

This policy is effectively a gradient ascent performed on a utility function (which ensures coverage uniformity), which interestingly can be derived from a global objective. This again shows the power of local utility design (compare this to the mechanism design for area exploration task).

In contrast to the area exploration, for the coverage problem, the swarm communication is not about one member announcing to the swarm that the one-time event of having found the target has occurred (where it is of importance to ensure this message is passed to the entire swarm as fast as possible) but about informing the (localized) swarm about each agent's state. For example, the coverage swarm should be robust against members dropping out, be it for natural reasons (battery depleted) or due to outside influence (something is interfering with the agent). The emphasis of the communication here is on ensuring that the relevant information is kept up to date at least locally, meaning that agents are aware of the state of their immediate neighbors. Consider an institutional rule in the swarm: *"broadcast your intent to vacate your current sector if you must leave it (e.g., due to low battery) so that neighbors can adjust"*.

This rule ensures robustness: if the task at hand is perimeter security then any patrol that cannot cover their shift must recruit others to help, ideally before this becomes an issue. If this is not possible (the agent fails without warning), or if communications are unreliable, the system can be made robust by well chosen defaults such as including in the agents' policies that they have a tendency to patrol slightly beyond their territory boundaries to create areas of overlap. If an agent is missing these areas will extend until the missing agent's territory is covered again. This is akin to the *checks and balances* in social systems where overlapping responsibilities can mitigate individual failures.

In this use case, an example of an emerging norm for the swarm is a timing convention: if agents slightly trespass on each other's territory, it is beneficial to do so when the other agent is not there (as the performance is calculated over the "time since the last visit" to a location it is beneficial to stagger the agent's arrival in any location). Agents that meet frequently could follow a simple protocol to de-synchronize to avoid future collisions, effectively falling into an out-of-phase patrolling rhythm. This simple process could be implemented through a single simple rule: *"if you encounter another agent regularly, slightly adjust your pace or route timing"*, much like the way fireflies synchronize (Strogatz, 2003) over time.

From a mechanism design viewpoint, the area coverage problem can be seen as a repetition of the area explore task, with repeated allocation of task to visit a specific location over time. This can be designed as a distributed optimization task: each agent works towards minimizing the time (for any location in its vicinity) that this location has not been visited. This local goal aligns with the global objective of minimizing worst-case coverage lag. Distributed patrolling algorithms in the literature are frequently based on cyclic paths or graph traversal. These algorithms enable near-optimal coverage, achieved with only local coordination, by designing each robot to react and adapt to their peers' presence and to share the workload with them.

In practice, implementing these ideas could look as follows:

- Agents drop virtual markers with timestamps as they move. They follow a policy: *"prefer to move toward areas with older timestamps (not visited recently) and avoid areas with very fresh timestamps"*. This naturally causes some areas to fall out of favor (negative feedback: the more an area is visited, the less attractive it becomes for a while).
- If an agent does not detect any marker from a neighbor in a long time, the probability that this neighbor is inactive increases and the agent slowly enlarges their patrol region to compensate (institutional rule for robustness).

- Conversely, if agents start to overlap heavily (finding each other's markers frequently), this starts to feel crowded . A possible adaptation is to either relocate or re-time the patrol (a spontaneous reallocation of effort). This could be driven by a rule: when a threshold of encountering other agents is reached the agent shifts its focus to an altogether new region which is currently not covered well at all.

Through these local rules, the swarm can self-organize into an effective persistent surveillance pattern. In our studies (e.g., Saffre et al. 2025, 2024, 2022a, 2022b) we have found that decentralized patrolling can be implemented efficiently and, under the right conditions, can outperform pre-scripted deployment pattern. In addition, our work has shown that such approaches can be resilient to agent failure or changes in the environment.

3.3 Comparative insights

The exploration and coverage scenarios illustrate how different (possibly opposing) emergent behaviors can be brought about by tailored and engineered policies for the agents implemented within appropriate social mechanisms:

- Exploration benefited from policies enforcing spatial dispersion and rapid information propagation (to achieve a one-time convergence on a target), effectively balancing positive and negative feedback to avoid both, premature clustering / aggregation and wandering (exploring) too long.
- Coverage benefited from policies encouraging spatial partitioning and routine formation (for monitoring), relying heavily on negative feedback (to prevent redundancy) and slow but steady adaptation to changes (for robustness).

Note that the role of communication is crucial: when looking for a target, the *"I found something"* signal is of utmost importance, it must be shared with the swarm as fast as possible. However, the criticality of information changes when we compare this to the coverage task: its still useful and improves overall efficiency but the task can be accomplished even when inter-agent communication is (temporarily) denied. Patrolling can be performed exclusively based on local markers. This is akin to the difference between an emergency broadcast and an ongoing operation where roles have been assigned in advance: the former is crucial while the latter can continue despite a breakdown in communication. This shows us that designing institutional constraints or norms must be a context-aware exercise.

Similarly, the two use-cases differ significantly in how *reward* or *utility* are structured: for the exploration task the reward is extremely asymmetric as all agents practically work towards the singular goal of locating the target; on the other hand, the coverage tasks uses uniformly distributed rewards (with a temporal parameter) as visiting any location contributes to the overall performance. This is akin to the *winner takes all* stance in economics (the agent finding the target *wins*), compared to an ongoing service provision (equal payout for any location) for the coverage task. This should not be surprising as the exploration task will benefit from designing competitive mechanisms into the agents' policies, while for the coverage task we would prefer to design mechanisms driving fairness and equality (load balancing). Note that we anthropomorphise the agents to be selfish or altruistic but this is to illustrate the framework only, agents simply have a preference for one next move over another without scheming to acquire an advantage ove their peers.

Ultimately, both tasks reinforce the central thesis: by conceiving agent behaviors as policies and employing social science concepts (feedback, norms, etc.) in designing those policies, we can achieve sophisticated emergent behaviors.

4. Discussion

The above framework and the two tasks we used indicate the cross-disciplinary approach to swarm engineering we envision. We will now discuss the benefits taking a page from the social science book, as well as some limitations to this. We briefly return to the concept of phase transitions because of their use as switches between behaviors, discuss the importance of robustness and once again state our preference for decentralized coordination in general. This manuscript is intended to start conversations on the subjects, so please discuss.

4.1 Benefits of the social science analogy

We think that the social scientist's view on behaviors of groups holds many benefits for the upcoming swarm engineer. For starters, Social Science has a head start, having investigated this for more than a century, with prominent examples being e.g., (Bon, 1895) (Group Psychology), (Durkheim, 1895) (Sociology) and (Freud, 1921) (Psychology). Secondly, it provides tools and lenses through which to anticipate or understand emergent outcomes in a richer and more complex manner as it essentially provides a means to talk about aspects of engineering a swarm. For example, *unintended consequences* (undesired and unexpected emergent behaviors), which is a well-studied problem in policy interventions (Schelling, 1978). Sociology and Economics have studied the notion of tipping points and power-law distributions and often provided a mathematical foundation for them. The use of notions such as *a norm* or *an institution* might seem strange at first due to their perceived vagueness (which engineers loath) but they do have (can have) a concrete meaning in the area of multi-robot coordination.

While this might not give us the mathematical framework to formally design emergence, this vocabulary can help to systematically think (and brainstorm) about designing a specific swarm. In addition, this gives us the analytical tools such as game theory, network analysis and system dynamics modelling to study the policies (algorithms) we design. For example, we can model the partitioning of an environment into territories as a game, with equally sized territories representing an equilibrium. This model can then be used to prove optimality, identify other equilibria or determine an equilibrium's stability.

4.2 Limits and differences

Despite all the mentioned analogies, there are of course significant differences between human social systems and robotic swarms. Members of the former have diverse preferences, goals, and the ability to consciously change strategies or even the rules of the game; while in the latter they typically share identical programming and have no independent agenda. Human decision making (Ajzen, 1991) can be vastly more complex than that of simple robotic platforms. This means some complex social mechanisms (like bargaining, deception, or coalition formation) might not have direct relevance in a deterministic swarm context. However, as autonomy and learning are introduced (e.g., reinforcement learning agents that adapt their policies), swarms might exhibit more agent-level *"self-interest"* or divergent behaviors that need alignment. This would be akin to the concepts of *selfish routing* in networks or *free-riding* in teams.

Our framework is directly applicable when the swarm is homogeneous and cooperative (all agents work toward the same global goal), because then the mechanism design can be simplified. In heterogeneous or adversarial multi-agent systems, additional game-theoretic considerations would come into play beyond the scope of this work. Clearly the proposed framework is a start, and more work is needed.

Another difference is that engineered systems allow us to explicitly implement or enforce rules, while in societies many rules can only be incentivized. For example, we can hard-code collision avoidance, whereas in human crowds one can only encourage avoiding collisions via cultural norms or signage. This means our design space is actually larger (we have more direct control), but also that some emergent phenomena in societies (stemming from bounded rationality or non-compliance) might not occur in swarms unless deliberately introduced. Therefore, we must carefully choose which social mechanisms to emulate and which to bypass via direct engineering (hard-coding). Often, the utility of emergent coordination is greatest in applications where central control is impossible (or at least very costly), similar to social systems which rarely have central planners. This makes emergent order a necessity for both, the social systems as well as the robotic swarms.

4.3 Phase transitions and design margins

We highlighted the concept of phase transitions in swarm behavior. A swarm designer should account for critical thresholds in agent connectivity (communication), population, or other parameters. If, for instance, the communication of a swarm is increasingly reduced (e.g., through jamming) then there will a point at which the swarm *collapses* and stop functioning as a whole. A similar effect can be achieved by iteratively reducing the swarm size (shooting down drones). By studying the system we may be able to identify where these phase transitions are located in the parameter space and then engineer safety triggers that change the overall behavior (or even the task) of the swarm whenever we get close to those identified regions in the parameter space. Examples of possible responses are (a) the deployment of additional drones, (b) the switch to environmental markers or (c) decreasing the search area. In social sciences, these strategies echo how communities ensure robustness by slightly overprovisioning (a safety margin).

Social models of abrupt change (like Schelling's segregation or riot threshold models) also provide a warning because they show us how varying a control parameter can lead to sudden (and undesired) shifts. In swarm terms, that could mean a graceful degradation suddenly turning into a breakdown. Recognizing this, designers might include backstops like mode-switching: if the swarm detects it's fragmenting (e.g., no communication from many agents), agents could switch to a fallback behavior such as returning to a rendezvous point. This is analogous to a society having emergency protocols when normal coordination fails (such as *"in case of fire gather at the meeting point"*).

4.4 Robustness and adaptation

We know that decentralization offers great potential for robustness, and if the swarm is coordinated through local communication and kept in check by emerging norms then there is no singular point of failure. That already helps. This robustness can be observed in self-organizing social systems, such as markets or interest communities, both of which can easily bounce back from shits in the environment, their members or supply.

In a distributed system, failure of a single agent (or even of a few agents) does not automatically have to result in the failure of the swarm, which can adapt with minimal performance loss.

Robustness also comes from adaptation: agent policies might include heuristics to adjust to new conditions such as changing the speed (coverage) or verifying the alarm (exploration). In addition, agents can use machine learning approaches to finetune their policies, much like in human societies the members learn through repeated interaction with one another. The swarm could learn, for example, that in a specific environment distance x is the optimal distance to keep to your neighbors. This value might be extremely hard to determine analytically but may emerge easily through trial and error. Using such learned adaptation may aid swarm coordination and enable the swarm to be successfully deployed in more diverse environments.

4.5 Decentralized coordination principles

Finally, we reflect on classic principles of decentralized coordination that are reinforced by our approach. One principle is *stigmergy*: indirect coordination (using the environment as a shared memory). In coverage, leaving behind information for others proved pouseful; this principle has broad applicability in swarm robotics, from foraging to construction tasks (Bonabeau et al., 1999). Stigmergy reduces the need for direct contact.

Another principle is *consensus emergence*: sometimes a swarm needs to reach agreement on a single course (e.g., all agents aligning their direction or timing). Distributed consensus algorithms (like averaging, voter models, or majority rule) are well-studied in control theory (Olfati-Saber et al., 2007) and mirror processes of opinion dynamics in groups. Our framework's emphasis on norms is related: a norm is effectively a consensus on a pattern of behavior. Ensuring consensus in a swarm often requires certain connectivity conditions; this ties to the above made point about ensuring communication networks remain intact enough (or periodically connecting the swarm if fully connected networks are impossible, maybe through mobile agents that shuttle information).

A third principle is *division of labor*, which we saw emerging in both tasks. In natural swarms and social insects, division of labor can arise via response thresholds or simple specialization rules. We similarly achieved partitioning of space (each agent covering *"their"* part). Encouraging division of labor without rigid assignment increases flexibility: if one agent is removed, others will cover its role. This is desirable in military swarms especially, where losses are expected.

Overall, the social science-inspired framework has the potential to help with the designing of a swarm as well as the explaining of its actions and behaviors. It allows us to articulate why a swarm did something in familiar terms, bridging the gap between human understanding and machine behavior. As autonomous swarms become more prevalent, such explanations can help operators trust and effectively integrate these systems.

Acknowledgement: This work was supported in part by the European Defence Fund (EDF) 2023 project SWARM-C3 (Command, Control and Communications for Multi-X Swarms). Views and opinions expressed are however of the author(s) only and do not necessarily reflect those of the European Union or the European Defence Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

Ajzen, I., 1991. The theory of planned behavior. *Organizational Behavior and Human Decision Processes*, 50(2), 179-211. Theories of Cognitive Self-Regulation.

Axelrod, R., 1986. An Evolutionary Approach to Norms. *American Political Science Review*, 80(4), 1095–1111.

Beckers, R., Holland, O. E., Deneubourg, J.-L., 2000. *Fom Local Actions to Global Tasks: Stigmergy and Collective Robotics*. Springer Netherlands, Dordrecht, 1008–1022.

Bon, G. L., 1895. *The Crowd: A Study of the Popular Mind*. Fischer Unwin, London.

Bonabeau, E., Dorigo, M., Theraulaz, G., 1999. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, NY.

Burgard, W., Moors, M., Stachniss, C., Schneider, F. E., 2005. Coordinated multi-robot exploration. *IEEE Transactions on Robotics*, 21(3), 376–386.

Cortés, J., Martínez, S., Karataş, T., Bullo, F., 2004. Coverage control for mobile sensing networks. *IEEE Transactions on Robotics and Automation*, 20(2), 243–255.

Couzin, I. D., Krause, J., Franks, N. R., Levin, S. A., 2005. Effective leadership and decision-making in animal groups on the move. *Nature*, 433(7025), 513–516.

Durkheim, E., 1895. *The Rules of Sociological Method*. Felix Alcan, Paris. Translated later into English by W.D. Halls.

Dussutour, A., Fourcassié, V., Helbing, D., Deneubourg, J.-L., 2004. Optimal traffic organization in ants under crowded conditions. *Nature*, 428(6978), 70-73. https://doi.org/10.1038/nature02345.

Freud, S., 1921. *Group Psychology and the Analysis of the Ego*. International Psycho-Analytical Press, London. Available via Project Gutenberg; original German title: Massenpsychologie und Ich-Analyse.

Hildmann, H., Kovacs, E., Saffre, F., Isakovic, A. F., 2019. Nature-Inspired Drone Swarming for Real-Time Aerial Data-Collection Under Dynamic Operational Constraints. *Drones*, 3(3). https://www.mdpi.com/2504-446X/3/3/71.

Hildmann, H., Nicolas, S., Saffre, F., 2012. A Bio-Inspired Resource-Saving Approach to Dynamic Client-Server Association. *IEEE Intelligent Systems*, 27(6), 17-25.

Holland, J. H., 1995. *Hidden Order: How Adaptation Builds Complexity*. Addison-Wesley, Reading, MA.

Hurwicz, L., Reiter, S., 2006. *Designing Economic Mechanisms*. Cambridge University Press, Cambridge, UK.

Husain, Z., Al Zaabi, A., Hildmann, H., Saffre, F., Ruta, D., Isakovic, A. F., 2022. Search and Rescue in a Maze-like Environment with Ant and Dijkstra Algorithms. *Drones*, 6(10). https://www.mdpi.com/2504-446X/6/10/273.

Johnson, S., 2002. *Emergence: The Connected Lives of Ants, Brains, Cities and Software*. Penguin science : culture, Penguin.

Miller, J. H., Page, S. E., 2007. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton University Press, Princeton, NJ.

Minati, G., Pessa, E., 2006. *Collective Beings*. Contemporary Systems Thinking, Springer US.

Okabe, A., Boots, B., Sugihara, K., Chiu, S. N., 2000. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. 2nd edn, John Wiley & Sons.

Olfati-Saber, R., Fax, J. A., Murray, R. M., 2007. Consensus and cooperation in networked multi-agent systems. *Proceedings of the IEEE*, 95(1), 215–233.

Ostrom, E., 1990. *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press, New York, NY.

Saffre, F., Hildmann, H., Anttonen, A., 2023. Force-Based Self-Organizing MANET/FANET with a UAV Swarm. *Future Internet*, 15(9). https://www.mdpi.com/1999-5903/15/9/315.

Saffre, F., Hildmann, H., Karvonen, H., Lind, T., 2022a. Monitoring and Cordoning Wildfires with an Autonomous Swarm of Unmanned Aerial Vehicles. *Drones*, 6(10). https://www.mdpi.com/2504-446X/6/10/301.

Saffre, F., Hildmann, H., Karvonen, H., Lind, T., 2022b. Self-swarming for multi-robot systems deployed for situational awareness. T. Lipping, P. Linna, N. Narra (eds), *New Developments and Environmental Applications of Drones*, Springer International Publishing, Cham, 51–72.

Saffre, F., Karvonen, H., Hildmann, H., 2024. Wild swarms: Autonomous drones for environmental monitoring and protection. T. Westerlund, J. Peña Queralta (eds), *New Developments and Environmental Applications of Drones*, Springer Nature Switzerland, Cham, 1–32.

Saffre, F., Karvonen, H., Hildmann, H., 2025. The Why and How of Polymorphic Artificial Autonomous Swarms. *Drones*, 9(1). https://www.mdpi.com/2504-446X/9/1/53.

Schelling, T. C., 1971. Dynamic models of segregation. *Journal of Mathematical Sociology*, 1(2), 143–186.

Schelling, T. C., 1978. *Micromotives and Macrobehavior*. W. W. Norton & Company, New York, NY.

Strogatz, S. H., 2003. *Sync: The Emerging Science of Spontaneous Order*. Hyperion.

Sutton, R. S., Barto, A. G., 1998. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA.

Young, H. P., 1993. The Evolution of Conventions. *Econometrica*, 61(1), 57–84.