



WHEN DRONES LEARN FROM EACH OTHER: **SOCIAL LEARNING IN AERIAL SWARMS**

■ Marco Dorigo and Mary Katherine Heinrich – DOI: <https://doi.org/10.1051/ejn/2026310>

■ IRIDIA, Université libre de Bruxelles – mdorigo@ulb.ac.be – mary.katherine.heinrich@ulb.be

Imagine a fleet of small drones exploring a disaster zone. No one tells them what to do. They have no prior knowledge of the site. Yet, within minutes, they have mapped the area, located the survivors, and divided the tasks among themselves—without ever connecting to a central computer. How? They learned from each other, on the fly.

This image is, of course, still idealised. Today's drone swarms are not yet capable of such fluid collective adaptation in the wild. But the direction is clear. As drones become cheaper, lighter, and more capable, the focus is shifting from the autonomy of individual drones to the collective behaviour of drone swarms. The issue is no longer only how to make a drone fly autonomously. It is how to make many drones operate together in a way that remains adaptive, scalable, and robust when conditions are uncertain and changing [1, 2].

For most of the history of robotics, a robot was essentially a solitary machine. Even when several robots were deployed together, each typically executed its own plan. The idea that robots might benefit from observing one another, exploiting each other's discoveries and adjusting their behaviours based on what others had learned, was more a speculative possibility than a practical design

principle. The solitary view works well as long as the robot is the natural unit of action. It becomes less convincing when the task itself calls for many robots operating together.

That is where aerial swarms become useful. A swarm of drones can do things that a single drone cannot do well: cover large areas quickly, inspect expansive infrastructure, monitor dynamic environments, or search in parallel. But once we move from one drone to many, coordination becomes the central problem. Defining in advance how each drone should behave becomes increasingly difficult as the swarm grows and the environment becomes less predictable. A central controller may simplify coordination, but it also creates a single point of failure, scales poorly, and often requires a global view of the system that can be difficult to obtain.

Social learning offers one possible way of addressing this problem. In biology, social learning [3, 4] is a

process by which an individual emulates others rather than discovering everything alone. Birds can learn routes from one another. Fish can emulate informed neighbours. Social insects exploit signals produced by nestmates or traces left in the environment. Humans, of course, rely on social learning on a massive scale [5].

For a swarm of drones, the same general idea applies. When a drone observes the action of another drone, it might copy that action directly, or imitate a larger behaviour, strategy, or discovery inferred from that action. If one drone discovers a safe passage, a promising search area, a landing site, or a region of interest, the rest of the swarm should be able to benefit from that discovery quickly, or the collective wastes time. If emulation spreads effectively through local observations, the swarm could become more efficient, more adaptive, and more robust.


What matters is that drones can benefit from each other's experience without losing the diversity and autonomy needed to keep exploring. The challenge is therefore not simply to learn socially by observing others, but to do so while preserving independent action and independent learning.

Where physics enters the picture

At first glance, social learning in a swarm of drones may look like a topic belonging mainly to biology, robotics, or artificial intelligence. But the connection with physics becomes clearer once we ask a simple question: how does useful information spread through a group of moving agents under real physical constraints? This question matters because the main advantages of a swarm [6, 7]—covering large areas, distributing tasks, and remaining functional even if some drones fail—depend on whether drones can exploit each other's discoveries without relying on perfect communication or central coordination. In a swarm of drones, communication is not free, sensing is not perfect, movement continuously reshapes the interaction network, and energy is always limited. For these reasons, social learning in drone swarms is not just a matter of designing good decision processes or interaction rules. It is also a matter of understanding how information propagates in the time-varying interaction network of an embodied dynamical system.

A simple example makes the issue concrete. Imagine a swarm searching inside a partially collapsed building. One drone discovers a traversable, relatively safe corridor. How can this discovery be shared with the rest of the swarm?

Consider three illustrative possibilities. The first route is direct instruction: the drone explicitly informs others nearby of its discovery. The second is emulation based on observation: nearby drones infer that a useful route has been found, and bias their own actions or goals accordingly. The third is coordination through a shared representation of the environment, such as a temporary




▲ FIG. 1: Three ways that information can spread in a drone swarm. Direct communication (left) is fast but might saturate the network when many drones send messages simultaneously. Observation (centre) is cheap but often ambiguous: nearby drones' states and behaviours must be inferred from sensor information. Shared environmental traces or distributed maps (right) enable efficient coordination but require safeguards against unrepresentative or outdated information that could mislead the group. The choice of mechanism depends on the task and on physical conditions.


distributed map or a digital pheromone-like field marking promising regions.

These are three ways information can spread through the swarm, but they do not share the same properties (see Fig. 1). Direct communication can be effective, but it may saturate the network when many drones need to exchange messages. Observation is cheaper but usually more ambiguous. Shared maps or digital traces can be a powerful tool for asynchronous coordination, but they require some way of preventing unrepresentative or outdated information from misleading the group. The engineering question is therefore not which communication strategy is universally best. The right choice depends on the task and on the operational conditions.

The key issue for social learning is then one of balance. A swarm that emulates too little might fail to exploit collective discoveries. A swarm that emulates too much might converge too quickly on a poor option or might become fragile because all members follow the same misleading cue. Social learning is therefore not simply a matter of increasing connectivity or imitation. ●●●

▼ FIG. 2: The interaction trade-off: from incoherence to rigidity. The interaction trade-off in drone swarms, and its relevance to social learning. With too little interaction (left), useful discoveries remain confined to individual drones and the swarm fails to benefit collectively. With too much interaction (right), the swarm converges rapidly but becomes rigid: a single erroneous cue can propagate to all members and dominate collective behaviour. The design challenge is to operate in the intermediate regime (centre), where information spreads efficiently enough to support effective coordination while preserving the autonomy and diversity that allow the swarm to remain robust under changing conditions.






▲ **FIG. 3:** From individual intelligence to collective intelligence. In a swarm, intelligence is not concentrated in a single unit. It resides partly in each drone's individual capabilities (sensing, flight control, local decision making, and uncertainty estimation), partly in the interactions among drones (observing and exchanging information with nearby drones and imitation with uncertainty), and partly in the collective dynamics that emerge from those interactions (emergent behaviour, shared knowledge, and adaptive coordination). This layered structure means that the swarm can be more capable than any of its members, without requiring central control.

It is a matter of finding the right balance between independence and emulation.

Physicists are used to systems in which too little interaction between elements prevents coherent collective behaviour, while too much interaction suppresses flexibility. Something similar happens here (Fig. 2). With too much interaction between individuals, the swarm becomes overly rigid: it might converge quickly but also become brittle and overreact to poor cues. In social learning, the challenge is to achieve enough emulation for useful behaviours to spread, while preserving enough flexibility for the swarm to remain adaptive when conditions

▼ **FIG. 4:** The connection between physics and social learning in drone swarms. Core challenges of swarm social learning (right) relate to well-developed areas of physics (left). How a discovery propagates through a drone swarm is determined in part by how information spreads in a dynamic interaction network, in which the topology changes as drones move. A drone deciding whether to trust a social signal includes problems of noise, measurement error, and uncertainty propagation: the same concepts used to analyse imperfect sensors and stochastic processes. The exploitation–exploration trade-off of an individual drone in a swarm is closely related to the interaction problem underlying coherence in many-body systems. All these challenges are further shaped by physical constraints—energy budgets, communication bandwidth, and flight kinematics—which are inherent to the system rather than externally imposed. Physics does not resolve these challenges in drone swarms, but it provides concepts and formalisms to address them precisely.



Physics provides key concepts; the swarm is the target system

change. This matters because conditions in the field are rarely static: wind shifts, GPS availability varies, obstacles appear unexpectedly, and air and light conditions can alter visibility. In such circumstances, social learning must remain selective and able to adapt. An individual should not only be able to exploit what others have found; it should also retain the capacity to discount information that is unreliable and to learn independently.

This also means that not all observations should carry the same weight. The swarm should learn not only based on members' observations, but also on their reliability. This, too, reflects a physical aspect of the problem. Measurements are noisy, estimates propagate through the swarm, and errors can accumulate, cancel out, or amplify. If a swarm is designed as though all signals were equally reliable, it will often behave poorly. If uncertainty is treated as part of the observation, the collective can become much more robust.

Seen in this way, the issue is not only how individual drones emulate and observe, but also how the swarm as a whole combines partial and imperfect signals. Intelligence in such a system does not reside in one place. Part of it resides in each drone, part in the interactions among drones, and part in the collective dynamics that emerge from those interactions (Fig. 3).

Shifting the view from individual to collective intelligence also helps explain why comparisons with animal or human societies are so appealing [6, 8]. Biological inspiration can be useful, but it should not be taken too literally. Drones are not birds: biological metaphors can be valuable, but provided they are translated into algorithms with care.

Biology is useful in showing that robust social learning is possible under severe constraints, including partial information, delayed communication, noisy sensing, and limited memory. These are precisely the kinds of constraints that drone swarms face.

Why it matters now

Drone swarms are slowly moving from the laboratory toward operation in the wild and real applications [9–12]: inspection, environmental monitoring, agriculture, infrastructure maintenance, and emergency response. In all these domains, the challenge is not merely to make drones fly. It is to make groups of autonomous drones adapt in the field, under uncertainty, without requiring ideal communication conditions.

Social learning becomes attractive in these circumstances, but it is not a complete answer. It does not eliminate the need for reliable sensing, robust flight control, or careful human oversight. Nor does it guarantee desired collective outcomes: the same mechanisms that allow useful behaviours to spread can also allow mistakes to spread. The challenge is therefore to design social learning mechanisms so that the swarm can benefit from shared experience without opening up new vulnerabilities.

This is what makes the subject scientifically rich. At its core, social learning in drone swarms raises a deceptively simple question: how does useful information become a collective resource in a moving, noisy, energy-limited physical system? Once phrased in this way, the question clearly goes beyond drones.

The more general issue, increasingly important in science and engineering, is: under what conditions can a group become more capable than its members, not because of central control, but because of the interactions among them? Physics does not answer that question by itself. But it provides much of the language, many of the concepts, and several of the tools needed to address it well (Fig. 4). ■

About the Authors



Marco Dorigo is the director of IRIDIA, the artificial intelligence laboratory of the Université libre de Bruxelles, where his work focuses on swarm intelligence and robotics, and a Research Director at the Fund for Scientific Research – FNRS (F.R.S.-FNRS), Belgium.



Mary Katherine Heinrich is a senior researcher at IRIDIA, the artificial intelligence laboratory of the Université libre de Bruxelles, and a Research Associate at the Fund for Scientific Research – FNRS (F.R.S.-FNRS), Belgium. Her research focuses on swarm robotics.

References

- [1] M. Dorigo, G. Theraulaz, and V. Trianni, *Science Robotics* **5**(49): abe4385 (2020)
- [2] M. Dorigo, G. Theraulaz and V. Trianni, *Proceedings of the IEEE* **109**(7), 1152 (2021)
- [3] E. Leadbeater and L. Chittka, *Current biology* **17**(16), R703 (2007)
- [4] C. Grueter and E. Leadbeater, *Trends in ecology & evolution* **29**(3), 177 (2014)
- [5] W. Hoppitt and K. N. Lala, *Social learning: an introduction to mechanisms, methods, and models*, Princeton University Press (2013)
- [6] E. Bonabeau, M. Dorigo and G. Theraulaz, *Swarm intelligence: from natural to artificial systems*, Oxford University Press (1999)
- [7] H. Hamann, *Swarm robotics: A formal approach*, Vol. 221, Springer (2018)
- [8] N. Coucke, M. K. Heinrich, A. Cleeremans and M. Dorigo, *Frontiers in Robotics and AI* **10**, 1030416 (2023)
- [9] G. Vászrhelyi, C. Viragh, G. Somorjai, T. Nepusz, A. E. Eiben and T. Vicsek, *Science Robotics* **3**(20), eaat3536 (2018)
- [10] K. N. McGuire, C. De Wagter, K. Tuyls, H. J. Kappen and G. C. H. E. de Croon, *Science Robotics* **4**(35), eaaw9710 (2019)
- [11] X. Zhou, X. Wen, Z. Wang, Y. Gao, H. Li, Q. Wang, T. Yang, H. Lu, Y. Cao and C. Xu, *Science Robotics* **7**(66), eabm5954 (2022)
- [12] W. Zhu, S. Oguz, M. K. Heinrich, M. Allwright, M. Wahby, A. L. Christensen, E. Garone and M. Dorigo, *Science Robotics* **9**(96), ead15161 (2024)

Challenges in Vacuum Measurement for Thermomechanical Material Testing

Thermomechanical analysis (TMA) relies on precise environmental control – especially vacuum conditions – to ensure reliable results. Even at low pressures, residual gases can affect measurements through adsorption, oxidation, and altered heat transfer, which is critical in high-precision and cryogenic applications.



A core challenge is accurate pressure measurement across a wide range—from rough vacuum ($\sim 10^2$ mbar) during evacuation to high vacuum ($\leq 10^{-6}$ mbar) during testing. No single sensor principle can cover this range with sufficient accuracy.

Modern systems therefore combine technologies: Pirani sensors for fine vacuum and cold cathode sensors for high vacuum. Thyracont addresses this with two approaches: the VSP series (Pirani) for stable fine vacuum measurement, and the VSM series, which integrates Pirani and cold cathode technologies for continuous monitoring into the high vacuum range.

Controlled activation of the cold cathode reduces wear and ensures long-term stability while maintaining continuous measurement—essential for reproducible TMA results.

As thermomechanical testing becomes increasingly sensitive, vacuum measurement is increasingly recognized as a key parameter for ensuring experimental accuracy and reproducibility. ■

For more information, contact:

Thyracont Vacuum Instruments
<https://thyracont-vacuum.com/en>
 Phone: +49 851 95986-0