Emergent Naming System in an Unstructured Environment: a Shortest-Path Discovery Case Study

Nicolas Cambier[1], A.E. Eiben[1], and Eliseo Ferrante[1,2]
[1]Computational Intelligence Group, Vrije Universiteit Amsterdam, The Netherlands
[2]Technology Innovation Institute, Abu Dhabi, United Arab Emirates
n.p.a.cambier@vu.nl

Abstract

It is often postulated that robots will eventually face conditions, whether on extraterrestrial bodies or deep underwater, that could not have been predicted by their designers. In such conditions, truly autonomous robots should be able to describe and talk about their environments in order to collectively find appropriate solutions. We designed an emergent naming systems for such purposes.

This paper focuses on a shortest-path discovery scenario in an unstructured environment, where landmarks are collectively named, by a swarm of robots, as they are discovered. The robots use those landmarks as beacons for navigation and score them according to their relevance to the task at hand. Meanwhile the naming system enables the swarm to update these scores asynchronously, using very little bandwidth.

We compare our naming-based navigation performances with swarms that do not communicate and swarms with prior knowledge of the environment, and find that our approach performs similarly to the latter. This has significant implications on the link between space conceptualisation and language, as this proto-language enables the robots to find a topological path without individually mapping the environment.

Introduction

Localisation is of paramount importance for achieving collective behaviours in a decentralised manner, as agents need to navigate between sites and to share their opinions thereof. Tasks that benefit from a shared reference frame include collective decision-making (Reina et al., 2018) (i.e. selecting the best option between two or more alternatives), foraging (Miletitch et al., 2018) (which can serve for, e.g., search-and-rescue missions), and area coverage and monitoring (Parrott et al., 2020; Tinoco and Oliveira, 2022). However, real environments such as caves, rarely provide the infrastructures required for global positioning. In particular, unstructured environments can dramatically reduce the reliability of such systems as global navigation satellite systems (GNSS), and odometry (i.e., using past control commands to estimate one's current position relative to a known starting point).

Swarm robotics, which takes inspiration from the behaviour of social insects, provides many alternatives to position-based approaches. Unfortunately, such approaches depend on highly stochastic navigation (Correll and Martinoli, 2011; Firat et al., 2018; Nouyan et al., 2009), and are therefore unsuitable for time-critical applications. Alternatively, stigmergy-inspired approaches (e.g. pheromones) either demand a high density of agents, who maintain a virtual gradient within their memory and update each other accordingly (Szymanski et al., 2006; Schmickl et al., 2007; Schmickl and Crailsheim, 2008), or custom actuator/sensor systems (Mayet et al., 2010; Khaliq et al., 2014; Fujisawa et al., 2014; Brenes-Torres et al., 2022) that are unpractical in the context of real world deployment. More reliable and economic alternatives exist (de Oca et al., 2011; Valentini et al., 2015; Reina et al., 2015; Gutiérrez et al., 2010), but they usually require distinct identifiers, shared by all robots, for each site of interest, which supposes prior knowledge and abilities that are unrealistic for most real world applications.

We propose to use naturally-occurring landmarks in unstructured environments, such as trees, rocks, and natural or artificial constructs, as beacons for navigation and, crucially, collective decision-making. Our approach does not require prior knowledge as these landmarks are discovered over the course of the experiment. Methods that enable robots to distinguish different objects or contexts from each other as they are discovered exist (Tani and Nolfi, 1999; Heinerman et al., 2017; Pitonakova, 2020), but learned models are often too complex for inter-robots communication (Kira, 2010). For this reason, this study presupposes the ability to individually distinguish landmarks, and focuses on the problem of self-organising a naming system for them. Such an ability is essential to social learning and cognition (Whiten et al., 2022; Coucke et al., 2023) and, therefore, to the deployment of robotic swarms in unknown environments (Cambier et al., 2020).

Approaches to enabling environment-driven communication can be separated into two categories: evolutionary approaches and cultural approaches. In the former, controllers are evolved, on a specific task, for robots equipped with communication devices, such as LEDs. Without it being required by the fitness function, the robots develop some form of signalling (Floreano et al., 2007; Ampatzis et al., 2008). As these solutions are evolved, they are highly specific to a given task and environment and therefore lack flexibility. Inversely, in cultural approaches, a basic language evolves over the course of the experiment through interactions between robots (Cambier et al., 2017, 2021; Miletitch et al., 2022). These approaches are based on language games; a model of the evolution of natural languages (Steels, 2011) and, more generally, of cultural consensus (Baronchelli et al., 2006). However, existing approaches either can not maintain a complete description of the represented landscape (Cambier et al., 2017, 2021) (as learned words are forgotten as soon as the robot's context changes), or also require a prior identification protocol for the salient features, such as using their position (Miletitch et al., 2022)–even if the latter are *a priori* unknown–, thus making the naming system somewhat redundant. In both cases, the emerged naming system is therefore useless to the task itself, although language consensus (i.e. knowing the same word, regardless of its meaning) can influence the robots' collective behaviour (Cambier et al., 2017, 2021, 2022). Despite past efforts (Cambier et al., 2017; Thenius et al., 2018; Miletitch et al., 2022), creating a useful "robotic culture" thus remains an open problem.

In this paper, we design an embodied landmark naming system, using language games, that is used for, and evolves in parallel with, the accomplishment of the task at hand. We test our approach on a collective navigation task, specifically a shortest-path discovery scenario, where robots must bring items from a source to their nest. Our method takes inspiration from stigmergic approaches, such as chemotaxis (Garnier et al., 2007; Parrott et al., 2020; Tinoco and Oliveira, 2022), e.g. ants pheromones, and trophallaxis (Schmickl et al., 2007; Schmickl and Crailsheim, 2008), whereby sharing and consuming food makes it sparser further away from the food source. Contrary to previous examples, our method does not require a high quantity of agents (Nouyan et al., 2009; Schmickl and Crailsheim, 2008; Parrott et al., 2020) nor custom sensors (Mayet et al., 2010; Khaliq et al., 2014; Fujisawa et al., 2014; Brenes-Torres et al., 2022), as robots use collectively named landmarks as "pheromone receptacles" and update each other asynchronously. We assess the efficacy of our approach in simulation, by comparing the navigation performances, with and without communication, and with and without prior knowledge.

Model & Method

Consider an environment with randomly distributed landmarks, wherein a swarm of robots is deployed. One of those landmarks is placed on a *source*, which is an area of interest (it might, e.g., contain resources, or require frequent monitoring). Another landmark is placed on the *nest*, where robots are initially deployed and where they need to return after visiting the source (e.g. to deposit resources or to recharge). We assume that any individual robot is able to differentiate the landmarks from each other (although the way this differentiation is performed might differ from one robot to another), but it does not need to be able to recognise the nest or the source from afar. Moreover, each landmark is presumed to be visible by a robot from at least one other landmark. In other words, the environment should be describable as a topological map with a connected network, where the nodes are landmarks and edges connect them if, and only if, the distance between the two landmarks is shorter than the robots' detection range. The swarm therefore has to find the shortest path between two *a priori* unknown nest and source, by topological navigation. For convenience, in this work, we used LEDs of different colours for each landmark, which the robots are able to detect using an on-board camera.

In the following, we first describe our navigation system and its use of virtual stigmergic cues to find the shortest path. Then, we explain how the communication required by the navigation can be used to simultaneously name the landmarks. Finally, we detail our experimental setup.

Collective Navigation

At a low-level, our robots are programmed with two basic behaviours that both use a target vector, which the robot progressively align their heading with. The first, and most im-

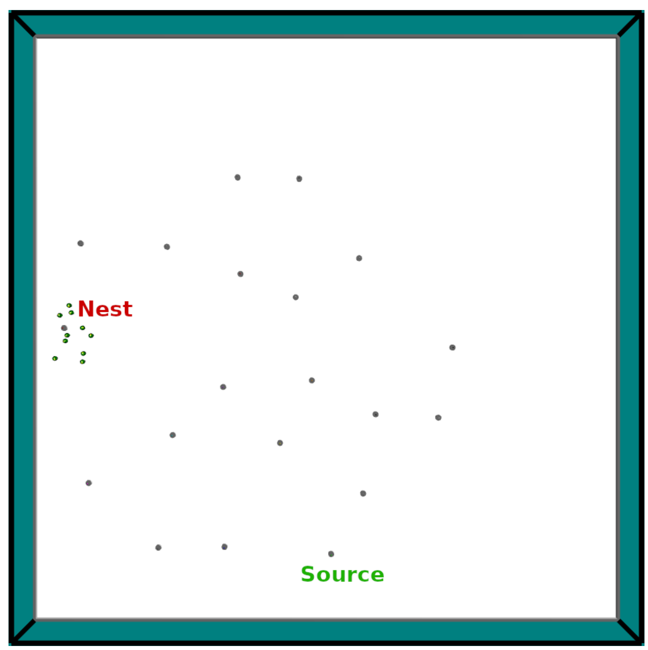

Figure 1: Example of a randomly generated environment. Nest and source landmarks are indicated.

portant, behaviour, $GOTO$, enables the robot to go towards a landmark. In that case, the target vector is simply the direction vector of said landmark, as detected by the robot's sensors. The second behaviour is a correlated random walk RW, whereby the target vector changes at a constant time interval RW_t. The new target vector is drawn from a uniform distribution centred around the previous target vector, within RW_θ degrees on either side. Both behaviours also implement an obstacle avoidance mechanism, that adjusts the target vector according to proximity sensor readings, in order to prevent robots from getting stuck.

The individual navigation behaviour of a given robot works as follows. First, the robot randomly selects a beacon landmark (according to rules detailed below), and uses its low-level behaviour $GOTO$ to drive towards it. On arriving, it starts a random walk, using RW, around the landmark. Importantly, it resumes $GOTO$ if it gets closer to another landmark than to the beacon, or if the beacon is beyond the range of its communication system. During RW, the robot can decide to leave the current landmark with probability p_{leave} and therefore selects and drives to a new beacon, as above. Obviously, this behaviour would result in a completely random path if the landmarks' selection probability was uniform, which is why we modulate this probability according to the swarm's past experiences, as follows.

Over the course of the experiment, robots use their sensors to detect the surrounding landmarks within their sensory range (including the nest landmark, which again, is indistinguishable from any other). For each detected landmark i that they do not yet know, they create a new element in their individual memory, which is characterised by its unique feature (in our case, the LED colour, although this could be a more complex model, such as a data stream cluster Heinerman et al. (2017)), and is associated to a base "pheromone level" $\tau_i = \tau^0$.

Let us note that τ is therefore associated to the landmarks, rather than to the path between them. From a graph theory point of view, this does not change our model's functioning, as we could create an equivalent graph with the landmarks as edges and the direct path between them as nodes (Dorigo et al., 2006). However, because of the high stochasticity of the environment, we predict that direct paths will largely outnumber the landmarks. Therefore associating τ to the landmarks reduces the memory requirements as well as simplifying information sharing (one word, instead of two with paths, uniquely identifies a point on the pheromone trail).

On reaching the beacon landmark i through $GOTO$, τ_i is updated by adding a parameter value ζ to it (τ_i is capped to 1). τ_i evaporates over time at a rate ρ (between 0 and 1) per second according to:

$$\tau_i' = (1 - \rho) * \tau_i \tag{1}$$

At this point, τ_i is a purely individual measure and would not suffice to find the shortest path (as demonstrated by our results). Therefore, while in RW around this beacon, a robot can broadcast the last l landmarks it visited, where l depends on the robot's bandwidth (we assume, for now, that landmarks have universal identifiers). All robots within communication range receive this update and increment each corresponding τ_i by ζ.

As the information offered by the broadcaster needs to be relevant to the navigation task, the probability Υ of sending this update depends on the time t_{visit} (in seconds) since this robot last visited an end landmark (i.e. the nest or the source). It is given by:

$$\Upsilon = (1 - \rho)^{t_{visit}} \tag{2}$$

i.e. the incentive to communicate reduces at the same rate as τ.

As explained above, at the start of an experiment or when deciding to leave a landmark, robots needs to pick a new beacon. This beacon is selected among the set L of landmarks detected in the current time-step[1], and that have not been visited since last visiting an end landmark (i.e. the nest or the source). Within L, the beacon is picked randomly with a probability weighted according to the square value of τ_i. In other words, landmarks with higher τ_i are more likely to be selected. Let us note that, if L is empty, the robot selects the least recently visited landmark, among those detected, as the beacon. At any point in time, if a robot can not see its beacon (due to obstacles or detection failures), it starts a random walk until the beacon becomes visible again.

Language Game

Robots can not always rely on universal identifiers for each salient feature in the environment. Indeed, this assumption supposes an unrealistic level of knowledge for most applications. Therefore, we designed a system to name landmarks as they are discovered, in a way that does not interfere with the communication required by the navigation task.

This system relies on the hearer-only Minimal Naming Game (MNG) (Baronchelli, 2016), a simplified version of language games that allows broadcasting, rather than peer-to-peer interactions, and is light on memory requirements (Trianni et al., 2016). In the MNG, each agent has a lexicon, i.e. a list of words, that is associated to a given topic. Agents can take either of two roles: speaker or hearer. In order to speak, the former selects a random word from its lexicon (it may have to create one first, if the lexicon is empty), and sends it. Upon reception, the hearer looks for that same word in its own lexicon. If that word is already present, the interaction is a success and the hearer forgets all other words. Otherwise, the interaction fails and the agent

[1] Let us note that some landmarks that are within detection range of the current landmark might not belong to L, as the robot might have driven too far in the opposite direction because of the random walk.

simply adds the word to its lexicon. With this process, the lexicons will first grow, as robots learn new words, before converging, quite suddenly, to a single words that uniquely identifies the topic (Baronchelli, 2016).

We established that, for the purpose of collective navigation, robots send and receive messages containing the identifiers of previously visited landmarks. We will now consider landmarks as topics in the MNG. In order to send its last few locations, a speaker can therefore select/create a random word associated to each given landmark, as in the MNG. On the receiving end, however, the received words need to be associated to a topic before the MNG can be played, as the corresponding landmark is unknown. As a matter of fact, identifying said landmark is the whole point of the interaction. Previous implementations of the MNG only worked for a single available topic, or when the topic could be clearly signalled before playing, e.g. through pointing. The guessing game is usually used in situations where, as in our case, the topic the speaker focuses on is ambiguous (Steels, 2011). Unfortunately, minimal versions of the guessing game do not exist, as feedback (i.e. confirming the intended topic after the game) is assumed to be essential for disambiguation. This would require both peer-to-peer interactions, and an ability, for both the speaker and the hearer, to indicate their intended topic after the fact. Nevertheless, those features are not necessary in our case, because our "minimal guessing game" is paired with a task.

We know, by our navigation model, that the first word of a message refers to the speaker's beacon, which the speaker remains closer to than to any other landmark, while the next few words' meanings are ambiguous. Therefore, the words received by the hearer can be separated into two categories: the first word, whose topic is known, and the other ones, whose topics are unknown. In the first case, the hearer can determine the topic simply by finding the landmark, among those detected in this time-step, that is closest to the relative position of the speaker robot (obtained through range-and-bearing communication). In the second case, the hearer looks over all its lexicons for the word, in order to list all the landmarks that already are associated to this word, and selects the topic randomly among those. In both cases, the MNG is eventually played with the selected topic only, and the corresponding landmark's pheromone level is incremented by ζ, as described in the previous subsection. In order to avoid homonyms (i.e. words that are identical but mean different things), the word is removed from any other lexicon that contains it.

Experimental Setup

Our experiments were performed in ARGoS (Pinciroli et al., 2012), a simulator with realistic physics, including collisions and actuation/sensory noise. Among the robot models available in ARGoS, we decided to implement our approach on the E-puck robots (Mondada et al., 2009). With its differential drive wheels, proximity sensors, range-and-bearing communication, and overhead $360°$ camera, this $75mm$-wide robot has all the requirements necessary for our approach.

Setup	Value	Description
Robot	$E-puck$	Robot Model
Runs	50	Runs per experimental setup
t_m	$5000s$	Experimental time
S	$10m$	Side length of the square arena
N	$\in \{10, 50\}$	Robots within the arena
msg_size	$\in \{0, 1, 2\}$	Message sizes in bytes
Q	20	Quantity of landmarks
$dist_S$	$6m$	Nest-Source distance
$dist_L$	$[1m, 3m]$	Min and max distance between landmarks
`Navigation`	Value	Description
`General`		
max_speed	10	Maximum speed in cms
cam_{range}	$3m$	Camera detection range
RAB_{range}	$1m$	Max communication range
`GOTO`		Beacon Navigation
p_{leave}	0.1	Probability to change beacon
`RW`		Correlated Random Walk
RW_t	$0.5s$	Heading change intervals
RW_θ	$54°$	Max heading change
`Pheromones`		
τ^0	0.1	Initial level
ζ	0.1	Increments
ρ	0.03	Evaporation rate

Table 1: Experimental parameters.

We defined the environment as a $10x10m$ square arena, populated with 20 unmovable cylinders, the landmarks, each topped with a LED of a unique colour. In order to normalise our observations, while assessing a wide range of environments, we distributed the landmarks according to the following protocol. First, the nest landmark is always placed in position $(0.5, 5)$ (i.e. on the left of the arena horizontally, and in the middle vertically). Then, the source landmark is placed at a distance of 6 metres from the nest, in a random direction. Finally, the 18 remaining landmarks are placed randomly and sequentially, with the condition that each is placed less than 3 metres away from at least one other landmark and no less than 1 metre away from any of them. A random environment, produced by this protocol, is presented in Figure 1.

In order to understand the impact of communication, and of the emergent naming specifically, we compared two variations of our approach. The first version, $Prior$, only implements the collective navigation, and assumes that the robots already know universal IDs for each landmark. This is our high-water mark. The second version, Minimal Guessing Game (MGG), implements both the navigation and the language game, as described above. Furthermore, we performed experiments with different bandwidths. As we set the size of words to one byte, we express this bandwidth, in number of bytes, as msg_size. We ran experiments with three values of $msg\_size = \{0, 1, 2\}$, which we selected because each of them makes a qualitative (rather than quantitative) difference. Indeed, $msg\_size = 0$ means that no communication occurs. As $Prior$ and MGG will perform individually (and, therefore, identically) without communication, $msg\_size = 0$ is the low-water mark. Then, $msg\_size = 1$ requires communication about the current beacon only, while $msg\_size = 2$ enables talking about the past, too. Meanwhile, from the point of view of MGG, moving from one to two bytes turns the naming from a straightforward MNG into a guessing game, as it now needs to disambiguate unknown topics.

Finally, we assessed our approach on two swarm sizes, 10 and 50. All our experimental parameters are given in Table 1. The navigation parameters were tuned manually for $Prior$, and then also applied to MGG.

Results

As our goal is to develop an on-line naming system with minimal impact on effectiveness, we compared the performances of MGG with swarms that do not communicate (as a baseline), and with swarms that have prior IDs for the landmarks ($Prior$), as a high-water mark.

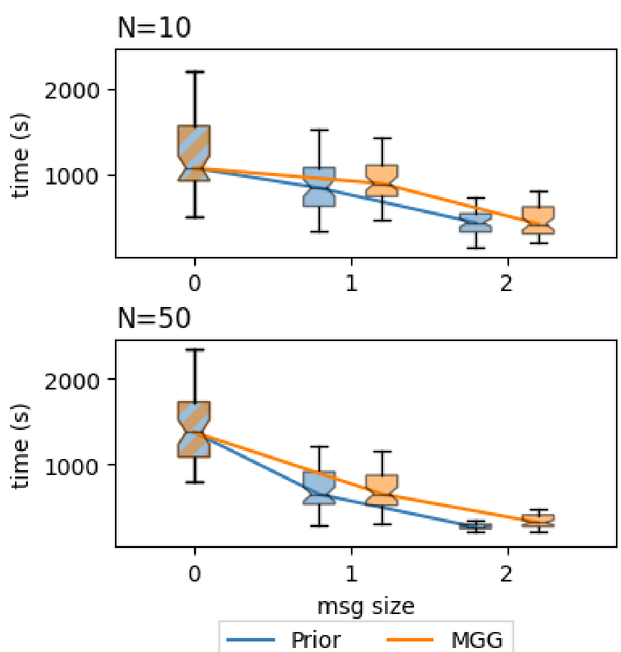

Figure 2: Average Lap Time at End of Experiments

N	ID	stat	p-value	Conclusion
10	*prior*	165	$5.1e-06$	Improves
	MGG	221	$5.8e-05$	Improves
50	*prior*	0	$7.6e-10$	Improves
	MGG	0	$7.6e-10$	Improves

Table 2: Effect, on lap time, of enabling communication ($msg\_size = 0 \rightarrow 1$, paired Wilcoxon test).

N	ID	stat	p-value	Conclusion
10	*prior*	40	$8e-09$	Improves
	MGG	18	$2.2e-09$	Improves
50	*prior*	12	$1.6e-09$	Improves
	MGG	36	$6.4e-09$	Improves

Table 3: Effect, on lap time, of enabling communication about undefined landmarks ($msg\_size = 1 \rightarrow 2$, paired Wilcoxon test).

Figure 2 shows the average lap times, for each configuration, at the end of the experiments. The lap time is the time it takes a robot to go from the nest to the source and back (repeated visits to either landmark outside of this order do not reset the counter). These results indicate that communication, as well as updating more landmarks at once, does improve the navigation performances. Indeed, the lap time reduces, and becomes more consistent, as msg_size increases. This observation is supported by Tables 2 and 3, which show the statistical significance of the improvements occurred by, respectively, enabling and expanding, communication. Additionally, the most notable difference between swarm sizes is observed when communication is disabled (lap time is much higher on average for $N = 50$), despite the behaviour being individual in that case, meaning that it should be identical across scales. We presume that this disparity is due to the increased physical interactions between robots within larger (and, therefore, denser) swarms, which elicits more avoidance and prevents robots from following straight paths.

The reason for the improvement caused by increasing msg_size is visible in Figure 3, which plots the pheromone levels on landmarks (with 50 robots), according to their distance to both the nest and the source. With this view, we see that landmarks that are close to the straight line between those two ends receive much more pheromones than the others. Incidentally, landmarks that are too close to the nest or to the source do not receive much pheromones as robots tend to ignore them in favour of a more direct path. Overall, smaller messages decrease the amount of pheromones that accumulate over landmarks on the optimal path, which reduces the difference between them and those that surround them. As a result, robots are more likely to explore and less likely to exploit the found path. As shown in Figure 4, this phenomenon, although less pronounced, is still

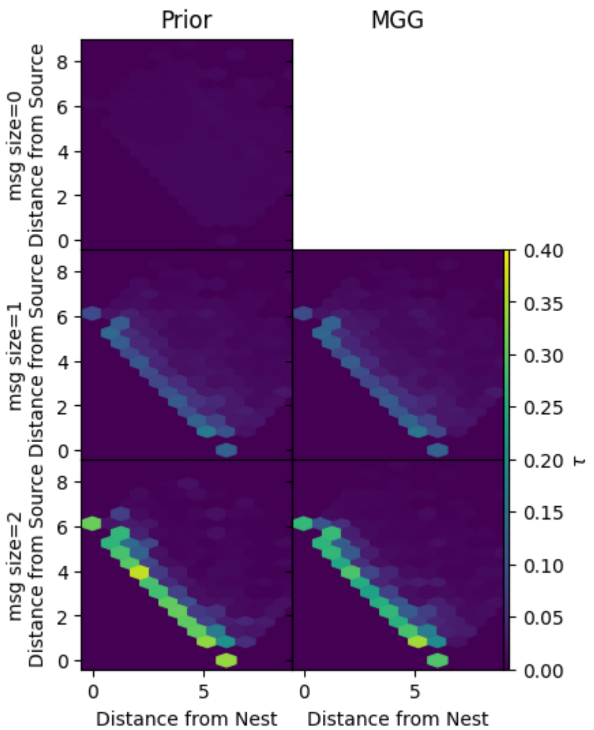

Figure 3: Pheromone levels on landmarks according to their distance from the nest and source, with 50 robots. Landmarks on the $y = 6 - x$ diagonal are on the Euclidean shortest path between the nest and the source.

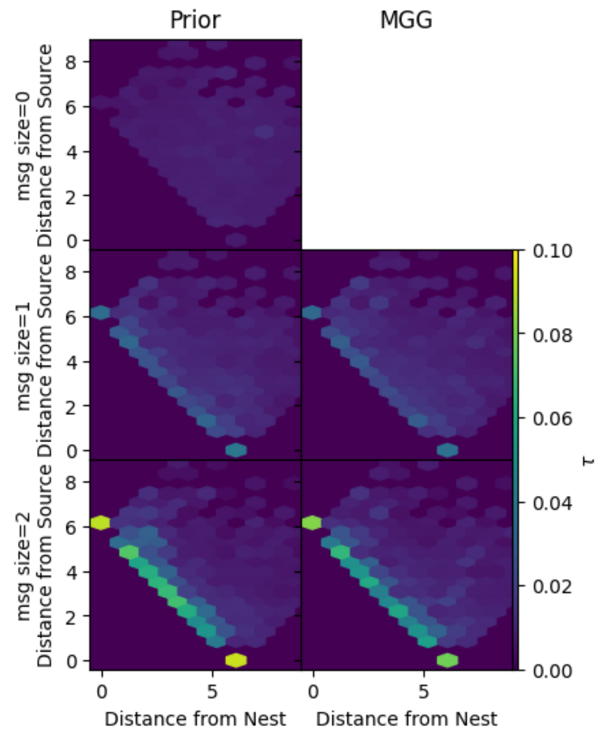

Figure 4: Pheromone levels on landmarks according to their distance from the nest and source, with 10 robots.

observed with a sparse swarm of 10 robots within a $100m^2$ arena. In other words, our approach enables efficient navigation in an environment whose area is more than two orders of difference larger than the area physically covered by the swarm, and over one order of difference with the maximum area covered by the latter's communication network.

The pheromones levels show that, without communication, the swarm does not find the shortest path, although lucky individuals might, which explains the high variability observed in Figure 2. This clearly demonstrates the benefit of cooperation in navigation. In contrast, there is no visible difference between $Prior$ and MGG. Therefore, we conclude that the latter does enable meaningful communication.

This conclusion is reinforced by looking at the emerged vocabulary. Figure 5 plots the evolution of the lexicon sizes over time, for four categories of landmarks: the nest, the source, the landmarks that are on the shortest topological path between those, and all the others. This figure shows that the vocabulary reduces to less than two words (on average) for the nest and shortest path landmarks quite fast. As the robots discover the source asynchronously, they initially tend to invent more words for that landmark, but eventually converge too. Conversely, other landmarks, as they are visited less frequently, never really converge vocabulary-wise, and a handful of synonyms for each of them is maintained over time. Overall, we observe that smaller values of either N or msg_size slightly slows down convergence on important landmarks, while higher values of both does the same for irrelevant landmarks. Finally, we observe a small amount of overlap (i.e. homonyms), as well as synonyms, despite MGG inhibiting them. However, their presence do not negatively impact the performance of our approach, meaning that there is no pressure for the robots to disambiguate those instances. Our hypothesis is that the swarm might create words that refer to several landmarks, but to the same concept. For example, it is possible that a word might apply to all the landmarks on the shortest path, or to landmarks that are too far away thereof, because updates on those landmarks are interchangeable.

N	msg_size	stat	p-value	Conclusion
10	1	501	0.19	Similar
	2	601	0.72	Similar
50	1	567	0.5	Similar
	2	195	$1.9e-05$	*MGG* is longer

Table 4: Statistical test on the similarity between $Prior$ and MGG w.r.t. lap time (paired Wilcoxon test).

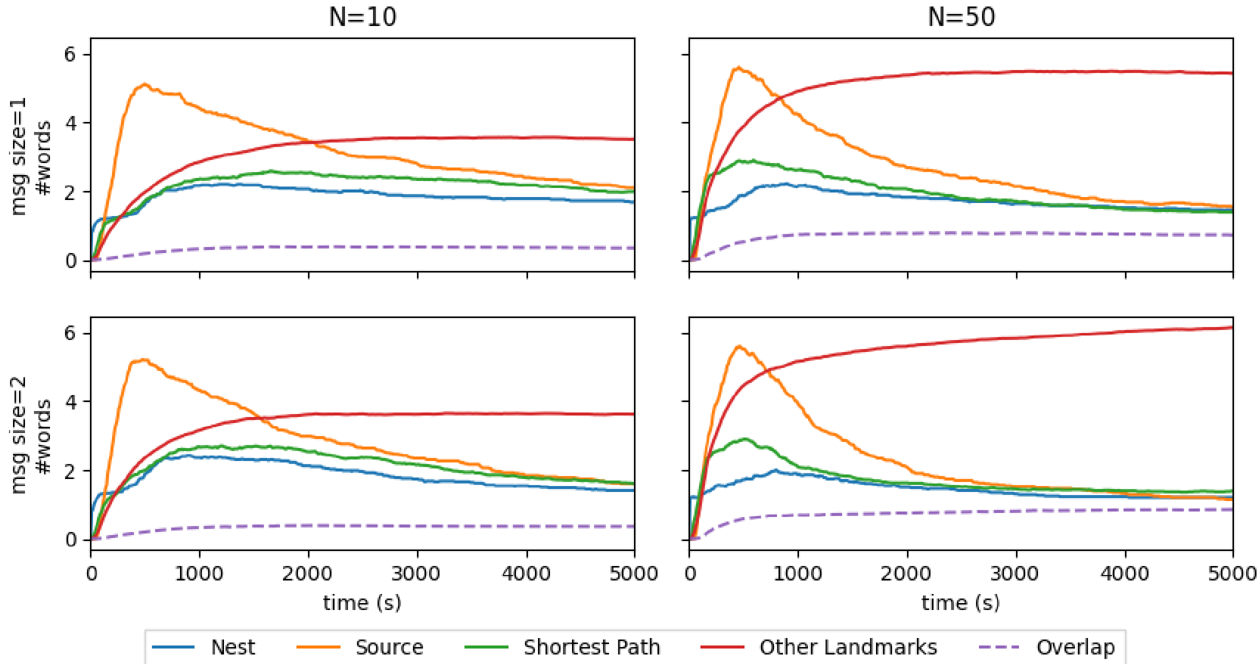

Figure 5: Average quantity of words associated to landmarks over time, when prior IDs are unavailable. *Nest* and *Source* represent a single landmark per run. *Shortest Path* and *Other Landmarks* represent sets of variable sizes, depending on the (random) landmark placement, and are therefore averaged per run before being averaged over the runs. Overlap is the average quantity of words that any landmark shares with any other.

To conclude, the swarm behaves similarly, whether it knows the landmark identifiers *a priori* or creates them on the fly. In order to gather more insights, we computed the statistical significance of this similarity. The results, presented in Table 4, indicate that $Prior$ is no more efficient than MGG, except for the case $N = 50$ and $msg\_size = 2$. Nevertheless, even in that case, MGG clearly improves performances compared to using a simple MNG (see Table 3 again). This means that MGG successfully disambiguates communication, even with a large quantity of potential topics. This is a very promising result as it suggests that prior knowledge of an environment may not be necessary.

Conclusion

In this paper, we introduced a localisation method for *a priori* unknown unstructured environments, whereby robotic swarms collectively name salient features, which they can then use as shared reference points for collective navigation and decision-making. We applied this method to a shortest-path discovery behaviour of our own design. Comparing our results, at two scales, to a high- and a low- watermark, we found that uninformed swarms performed very similarly to ones with prior information about the environment.

Our results provide new insights onto the link between social cognition, language and space conceptualisation. Indeed, humans tend to conceptualise space as a collection of salient features (Vasudevan et al., 2007), but not as a map (Kuipers, 2000). We showed that, through collective naming, agents can indeed learn to navigate within a large space, without individually mapping the topology of its salient features. Furthermore, we observed that synonyms are more frequently associated to concepts that are less relevant to the population, as the pressure to unambiguously refer to them is lessened. This is somewhat contrary to the popular "words for snow" theory (i.e. languages in cold climates would have more words referring to snow), which provides new data for an assumption that remains largely unfounded (Regier et al., 2016).

To the best of our knowledge, our approach is the first ever implementation of a robotic culture, which is flexible to environmental variations, useful for the task at hand, and does not assume a shared sensory space or detection model. The properties of language games make this culture adaptive as they are capable of taking new meanings into account (and actually emerged that way in the first place) and of including new agents, even if these agents already have a language (Steels, 2011). This approach is suitable for robotic systems with very limited bandwidth as messages of a couple of bytes suffice to drastically improve performances.

Our next step will be to assess the efficiency of our approach and its adaptability to environmental changes, such as adding/removing robots or landmarks, or moving the nest/source to another landmark, over the course of the experiment. This would clearly illustrate the benefits of naming landmarks on the fly, rather than relying on prior knowledge.

Moreover, we will further assess the scalability of our model with respect to the swarm size (although our results indicate that increasing this parameter improves performances, this tendency might reverse at very large scales), and its resilience to varying parameters on the environmental (e.g. quantity and distance between landmarks) and experimental (e.g. sensor ranges, navigation parameters, *etc*) levels. We are particularly interested in determining whether there is any setting whereby emergent naming outperforms prior identifiers.

Furthermore, we will work on extending our approach to more robotic applications. Area monitoring is readily available as we can just program robots to avoid, rather than seek, high pheromone levels (Parrott et al., 2020). Less straightforward adaptations will also be considered for path planning, rendez-vous, and collective decision-making in general. In that same spirit, we will implement our approach on actual robots, and couple it with existing algorithms to detect and distinguish the landmarks in more complex environments. A promising approach would be to directly link sensing data with a more complex language game as the latter enables agents to collectively categorise continuous signals (Steels, 2008; Baronchelli et al., 2010).

As we observed the existence of homonyms that do not affect the performance of the swarm, we formulated the hypothesis that some words had emerged, that referred to abstract concepts such as "the shortest path" or "irrelevant landmarks", rather than to individual landmarks. Such an emergent vocabulary would go beyond our intended design, which would be a promising feature for collective learning. We will therefore study the culturally evolved vocabularies in more details, in order to verify the emergence of abstract thinking in the swarm.

Acknowledgements

This work is supported by Technology Innovation Institute (TII), UAE.

References

Ampatzis, C., Tuci, E., Trianni, V., and Dorigo, M. (2008). Evolution of signaling in a multi-robot system: Categorization and communication. *Adaptive Behavior*, 16(1):5–26.

Baronchelli, A. (2016). A gentle introduction to the minimal naming game. *Belgian Journal of Linguistics*, 30(1):171–192.

Baronchelli, A., Felici, M., Loreto, V., Caglioti, E., and Steels, L. (2006). Sharp transition towards shared vocabularies in multiagent systems. *Journal of Statistical Mechanics: Theory and Experiment*, 2006(06):P06014.

Baronchelli, A., Gong, T., Puglisi, A., and Loreto, V. (2010). Modeling the emergence of universality in color naming patterns. *Proceedings of the National Academy of Sciences*, 107(6):2403–2407.

Brenes-Torres, J. C., Blanes, F., and Simo, J. (2022). Magnetic trails: A novel artificial pheromone for swarm robotics in outdoor environments. *Computation*, 10(6):98.

Cambier, N., Albani, D., Frémont, V., Trianni, V., and Ferrante, E. (2021). Cultural evolution of probabilistic aggregation in synthetic swarms. *Applied Soft Computing*, 113:108010.

Cambier, N., Frémont, V., and Ferrante, E. (2017). Group-size regulation in self-organised aggregation through the naming game. In *International Symposium on Swarm Behavior and Bio-Inspired Robotics (SWARM 2017)*, Kyoto, Japan.

Cambier, N., Miletitch, R., Burraco, A. B., and Raviv, L. (2022). Prosociality in swarm robotics: A model to study self-domestication and language evolution. In *Joint Conference on Language Evolution (JCoLE)*, pages 98–100. Joint Conference on Language Evolution (JCoLE).

Cambier, N., Miletitch, R., Frémont, V., Dorigo, M., Ferrante, E., and Trianni, V. (2020). Language evolution in swarm robotics: A perspective. *Frontiers in Robotics and AI*, 7:12.

Correll, N. and Martinoli, A. (2011). Modeling and designing self-organized aggregation in a swarm of miniature robots. *The International Journal of Robotics Research*, 30(5):615–626.

Coucke, N., Heinrich, M. K., Cleeremans, A., and Dorigo, M. (2023). Learning from humans to build social cognition among robots. *Frontiers in Robotics and AI*, 10:1030416–1030416.

de Oca, M. A. M., Ferrante, E., Scheidler, A., Pinciroli, C., Birattari, M., and Dorigo, M. (2011). Majority-rule opinion dynamics with differential latency: a mechanism for self-organized collective decision-making. *Swarm Intelligence*, 5(3-4):305–327.

Dorigo, M., Birattari, M., and Stutzle, T. (2006). Ant colony optimization. *IEEE computational intelligence magazine*, 1(4):28–39.

Firat, Z., Ferrante, E., Cambier, N., and Tuci, E. (2018). Self-organised aggregation in swarms of robots with informed robots. In *International Conference on Theory and Practice of Natural Computing*, pages 49–60. Springer.

Floreano, D., Mitri, S., Magnenat, S., and Keller, L. (2007). Evolutionary conditions for the emergence of communication in robots. *Current biology*, 17(6):514–519.

Fujisawa, R., Dobata, S., Sugawara, K., and Matsuno, F. (2014). Designing pheromone communication in swarm robotics: Group foraging behavior mediated by chemical substance. *Swarm Intelligence*, 8:227–246.

Garnier, S., Tache, F., Combe, M., Grimal, A., and Theraulaz, G. (2007). Alice in pheromone land: An experimental setup for the study of ant-like robots. In *2007 IEEE swarm intelligence symposium*, pages 37–44. IEEE.

Gutiérrez, A., Campo, A., Monasterio-Huelin, F., Magdalena, L., and Dorigo, M. (2010). Collective decision-making based on social odometry. *Neural Computing and Applications*, 19:807–823.

Heinerman, J., Haasdijk, E., and Eiben, A. (2017). Unsupervised identification and recognition of situations for high-dimensional sensori-motor streams. *Neurocomputing*, 262:90–107.

Khaliq, A. A., Di Rocco, M., and Saffiotti, A. (2014). Stigmergic algorithms for multiple minimalistic robots on an rfid floor. *Swarm Intelligence*, 8:199–225.

Kira, Z. (2010). Inter-robot transfer learning for perceptual classification. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*, pages 13–20.

Kuipers, B. (2000). The spatial semantic hierarchy. *Artificial intelligence*, 119(1-2):191–233.

Mayet, R., Roberz, J., Schmickl, T., and Crailsheim, K. (2010). Antbots: A feasible visual emulation of pheromone trails for swarm robots. In *ANTS Conference*, pages 84–94. Springer.

Miletitch, R., Dorigo, M., and Trianni, V. (2018). Balancing exploitation of renewable resources by a robot swarm. *Swarm Intelligence*, 12(4):307–326.

Miletitch, R., Reina, A., Dorigo, M., and Trianni, V. (2022). Emergent naming conventions in a foraging robot swarm. *Swarm Intelligence*, 16(3):211–232.

Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., Klaptocz, A., Magnenat, S., Zufferey, J.-C., Floreano, D., and Martinoli, A. (2009). The e-puck, a robot designed for education in engineering. In *Proceedings of the 9th conference on autonomous robot systems and competitions*, volume 1, pages 59–65. IPCB: Instituto Politécnico de Castelo Branco.

Nouyan, S., Groß, R., Bonani, M., Mondada, F., and Dorigo, M. (2009). Teamwork in self-organized robot colonies. *IEEE Transactions on Evolutionary Computation*, 13(4):695–711.

Parrott, C., Dodd, T. J., Boxall, J., and Horoshenkov, K. (2020). Simulation of the behavior of biologically-inspired swarm robots for the autonomous inspection of buried pipes. *Tunnelling and Underground Space Technology*, 101:103356.

Pinciroli, C., Trianni, V., O´Grady, R., Pini, G., Brutschy, A., Brambilla, M., Mathews, N., Ferrante, E., Di Caro, G., Ducatelle, F., et al. (2012). Argos: a modular, parallel, multi-engine simulator for multi-robot systems. *Swarm intelligence*, 6(4):271–295.

Pitonakova, L. (2020). Novelty detection with robots using the grow-when-required neural network.

Regier, T., Carstensen, A., and Kemp, C. (2016). Languages support efficient communication about the environment: Words for snow revisited. *PloS one*, 11(4):e0151138.

Reina, A., Bose, T., Trianni, V., and Marshall, J. A. (2018). Effects of spatiality on value-sensitive decisions made by robot swarms. In *Distributed Autonomous Robotic Systems: The 13th International Symposium*, pages 461–473. Springer.

Reina, A., Miletitch, R., Dorigo, M., and Trianni, V. (2015). A quantitative micro–macro link for collective decisions: the shortest path discovery/selection example. *Swarm Intelligence*, 9(2-3):75–102.

Schmickl, T. and Crailsheim, K. (2008). Trophallaxis within a robotic swarm: bio-inspired communication among robots in a swarm. *Autonomous Robots*, 25:171–188.

Schmickl, T., Möslinger, C., and Crailsheim, K. (2007). Collective perception in a robot swarm. In *Swarm Robotics: Second International Workshop, SAB 2006, Rome, Italy, September 30-October 1, 2006, Revised Selected Papers 2*, pages 144–157. Springer.

Steels, L. (2008). The symbol grounding problem has been solved. so what's next. *Symbols and embodiment: Debates on meaning and cognition*, pages 223–244.

Steels, L. (2011). Modeling the cultural evolution of language. *Physics of Life Reviews*, 8(4):339–356.

Szymanski, M., Breitling, T., Seyfried, J., and Wörn, H. (2006). Distributed shortest-path finding by a micro-robot swarm. In *Ant Colony Optimization and Swarm Intelligence: 5th International Workshop, ANTS 2006, Brussels, Belgium, September 4-7, 2006. Proceedings 5*, pages 404–411. Springer.

Tani, J. and Nolfi, S. (1999). Learning to perceive the world as articulated: an approach for hierarchical learning in sensory-motor systems. *Neural Networks*, 12(7-8):1131–1141.

Thenius, R., Moser, D., Varughese, J. C., Kernbach, S., Kuksin, I., Kernbach, O., Kuksina, E., Mišković, N., Bogdan, S., Petrović, T., et al. (2018). subcultron-cultural development as a tool in underwater robotics. In *Artificial Life and Intelligent Agents: Second International Symposium, ALIA 2016, Birmingham, UK, June 14-15, 2016, Revised Selected Papers 2*, pages 27–41. Springer.

Tinoco, C. R. and Oliveira, G. (2022). Pherocom: decentralised and asynchronous swarm robotics coordination based on virtual pheromone and vibroacoustic communication. *arXiv preprint arXiv:2202.13456*.

Trianni, V., De Simone, D., Reina, A., and Baronchelli, A. (2016). Emergence of consensus in a multi-robot network: from abstract models to empirical validation. *IEEE Robotics and Automation Letters*, 1(1):348–353.

Valentini, G., Ferrante, E., Hamann, H., and Dorigo, M. (2015). Collective decision with 100 kilobots: speed versus accuracy in binary discrimination problems. *Autonomous Agents and Multi-Agent Systems*, pages 1–28.

Vasudevan, S., Gächter, S., Nguyen, V., and Siegwart, R. (2007). Cognitive maps for mobile robots—an object based approach. *Robotics and Autonomous Systems*, 55(5):359–371.

Whiten, A., Biro, D., Bredeche, N., Garland, E. C., and Kirby, S. (2022). The emergence of collective knowledge and cumulative culture in animals, humans and machines.