

Enhanced Foraging in Robot Swarms Using Collective Lévy Walks

Johannes Nauta¹ and Stef van Havermaet and Pieter Simoens and Yara Khaluf

Abstract. A key aspect of foraging in robot swarms is optimizing the search efficiency when both the environment and target density are unknown. Hence, designing optimal exploration strategies is desirable. This paper proposes a novel approach that extends the individual Lévy walk to a collective one. To achieve this, we adjust the individual motion through applying an artificial potential field method originating from local communication. We demonstrate the effectiveness of the enhanced foraging by confirming that the collective trajectory follows a heavy-tailed distribution over a wide range of swarm sizes. Additionally, we study target search efficiency of the proposed algorithm in comparison with the individual Lévy walk for two different types of target distributions: homogeneous and heterogeneous. Our results highlight the advantages of the proposed approach for both target distributions, while increasing the scalability to large swarm sizes. Finally, we further extend the individual exploration algorithm by adapting the Lévy walk parameter α , altering the motion pattern based on a local estimation of the target density. This adaptive behavior is particularly useful when targets are distributed in patches.

1 Introduction

Coordination and self-organization in swarm robotics are key features that are inspired from natural social systems such as ant colonies [27], and enable a large group of robots to achieve collective efficiency that is not achieved by individual robots [29, 6, 12]. In robot swarms, each robot builds its knowledge of the world locally through its limited perception. Nevertheless, robots exploit direct and indirect communication to increase their effectiveness as a group by sharing information and make decisions conjointly. One of the well-studied examples in robot swarms is foraging, which is behavior commonly observed in social animals [4]. This task is intensively studied in robot swarms due to its importance as a metaphor for a large spectrum of robotic applications, including search and rescue, and resource exploitation (e.g., harvesting).

In foraging, the collective system searches for targets in an unknown environment. Such exploration attempts are bound to constraints, such as energy expenditure, and thus need to be optimized in order to ensure the survival of the swarm. Several studies have shown that when foraging considers individuals, random searches can be optimized depending on the distribution over targets [31, 23, 25, 38]. In these studies, random searches which follow Lévy walk patterns have been found to be more efficient than alternatives, such as the correlated random walk [3] and Brownian motion. In random walks, the trajectory of an individual is described by a sequence of flights, whose length and direction are chosen randomly. In the specific case

of Lévy walks, the flight lengths ℓ are sampled from a power-law distribution

$$p(\ell) \sim \ell^{-(\alpha+1)}, \quad (1)$$

where $0 < \alpha < 2$ is the Lévy parameter. In contrast to many exponentially decaying distributions, such as the normal distribution, power-law distributions have tails which are fat. This results in statistically relevant large values being sampled from the distribution, corresponding to execution of long flights in the case of a Lévy walk (for an extensive review, we refer the interested reader to [37]). Hence, Lévy walks alternate long bouts of straight line motion with Brownian-like motion, furthermore displaying scale-free behavior typical of power-laws. Exactly these long flights are responsible for the increase in the search efficiency needed for foraging animals to survive.

In more detail, random searches with Lévy parameter $\alpha \approx 1$ optimize the random search for an individual over a wide range of target distributions, such as homogeneous, sparse distributions [31, 3], heterogeneous, patchy distributions [34, 36, 35] and scale-free, fractal-like distributions [11]. Furthermore, Lévy walks are of interest due to the wide range of motion behavior that they encompass. For $\alpha \rightarrow 0$, we enter the ballistic regime where each individual only displays straight line motion. For $\alpha \geq 2$, the resulting motion is Brownian due to the power-law distribution converging to the normal distribution due to the central limit theorem. Intermediate values $0 < \alpha < 2$ show motion patterns in between both extremes, including an optimal value at $\alpha_{opt} \approx 1$ that interchanges long flight lengths with more Brownian-like behavior.

While individual foraging can be optimized by having the forager follow a Lévy walk, collective foraging is not optimized by simply having each individual follow a Lévy walk, due to physical limitations and finite size effects. In cases where the density of the swarm is high, collisions between individuals become increasingly likely, resulting in truncation of the current flight due to collision avoidance. Hence, the collective system as a whole might not display the characteristics typical for a Lévy flight, i.e. a power-law distribution over flight lengths. To the best of our knowledge, it is currently unclear if the resulting collective motion actually resembles a Lévy walk for an increasing swarm density. The purpose of this study is therefore twofold; (i) we first aim to understand the influence of the swarm density on the collective behavior of the random search, and (ii) provide an engineered approach to increase the search efficiency of the collective system. Specifically, we (i) analyze the distribution of the collective flight lengths, for which we show that this does not follow a Lévy characteristic power-law anymore above a certain (critical) swarm size. Afterwards, we (ii) develop a random walk strategy for each individual which ensures power-law distributions for larger

¹ Ghent University, Belgium, email: johannes.nauta@ugent.be

swarm sizes. We furthermore show that this strategy is able to increase the search efficiency, in two specific target distributions, which are common to observe in nature. This in turn benefits the collective foraging task in robot swarms.

2 State of the art

Within swarm robotics, random walks are inherently fundamental to a large number of applications. In tasks such as aggregation [26], where individuals cluster to form coherent groups or to detect particular environmental cues [15], initial search for other members of the swarm occurs using random walks. In collective exploration [32], for example for mapping unknown environments [14], individual motion determines the exploration efficiency. Furthermore, when communication is restricted within a small radius, such as in many collective decision making scenarios [30, 18] and collective foraging [4, 10], the local networks wherein information is transferred are highly dependent on the individual motion [16]. Remarkably, while the random walk constitutes a critical component of swarm robotics, research into the random walk characteristics of the individuals and the collective system is limited.

Moreover, while Lévy walks do appear as the optimal search strategy in collective foraging in sparse environments [28, 5, 9], these works do not consider large swarm sizes or high swarm densities. However, in dense swarms, the aforementioned collisions between individuals appear more frequent, resulting in truncation of the current flight due to the inherent collision avoidance of the individuals. Hence, this study aims to investigate the effect of the density on the characteristics of the swarm's collective random walk. We believe that our work can benefit collective foraging in swarm robotics, especially when large swarms are considered.

3 Model description

In the following, we discuss three building blocks of the swarm behavior model; (i) the environment in which the individuals and targets are embedded, (ii) the individual exploration model that describes the searching behavior applied by each robot, and finally (iii) the collective performance measures that we use to examine the efficiency of the proposed exploration approaches, in addition to the distribution of the emergent collective pattern.

3.1 Environment description

We study the flight length distribution by having a swarm of N robots move around in a two-dimensional $L \times L$ environment, where the boundaries are periodic. The periodic boundaries reflect an environment which is much larger than each individual robot and hence alleviates potential boundary effects, allowing us to study the influence of swarm density in isolation.

For measuring the search efficiency of the different random walk implementations, we distribute targets across the environment. In this work, we assume that targets are revisitable, which is analogous to *non-destructive* foraging in the literature (see e.g. [31]). Targets are defined by circles, each with radius r_t , and are detected by the robot upon entering the circle. This is equivalent to the robots searching for point-like targets with a detection radius of r_t , similar to previous implementations in related works [3, 36]. However, in contrast these works, we do not truncate the current flight whenever a target is detected. This difference is introduced because we empirically found truncation at target detection to result in the cluttering of robots

close to the targets, which does not correspond to random searches for targets. Studying behavioral differences at target detection is considered valid future research. In this study, we focus on two distinct target distributions:

Homogeneous: To study sparse, but homogeneous, densities, we uniformly distribute M targets over the environment. The number of targets needs to be controlled, since we aim to discuss sparse target distributions, for which we know the optimal Lévy parameter from the literature.

Heterogeneous: Many real-world environments are actually not homogeneous. Targets are often distributed in patches [19, 13, 33], where it is often assumed that the target distribution is only homogeneous within the patches. There are no targets outside of the patches. Whereas such environments vastly differ from the homogeneous case, it has been found that the optimality of the Lévy flight is robust to the different environments [34]. To study patchy environments, we distribute K_p non-overlapping, circular patches of radius $R < L$ over the environment and uniformly distribute M_p targets within each single patch.

3.2 Individual behavior

Next, we discuss the individual behavior of each robot within the swarm. As a baseline model, we combine Lévy walk exploration behavior with collision avoidance, and define this as the *individual Lévy walk* (ILW). This model is extended by enabling each robot to communicate relative positional information with its neighboring robots, while executing its own Lévy walk. The latest model is referred to as the *collective Lévy walk* (CLW). An additional extension, which specifically targets the maximization of the collective search efficiency is provided by adapting the Lévy parameter α , based on the predicted density of targets—in cases where targets are distributed heterogeneously. This model is referred to as the *adaptive collective Lévy walk* (ACLW). We present the three different individual models in the following sections. An overview of the (adaptive) collective Lévy walk is given in algorithm 1.

3.2.1 Individual Lévy walk strategy

Similar to [28], the individual Lévy walk, which includes collision avoidance, is executed by each individual. It is given as follows:

(a) each robot samples a flight length according to the power-law distribution of equation 1, where $\alpha = \alpha_{opt} = 1$. The power-law is truncated to ensure at least a minimum step size of $\ell_{min} = d_b$ is taken, in addition to an upper truncation of ℓ_{max} . Here, d_b is the physical diameter of the robot. The lower truncation ensures that steps smaller than the target detection radius $r_t = d_b$ are not taken into account since those are meaningless in a search context. The upper truncation is present to omit large flight lengths, which are unrealistic in real-world settings (e.g., larger than the dimensions of the environment). After sampling the flight length, the robot walks in a straight line with fixed velocity v .

(b) after traversing the sampled distance, each robot rotates for a number of steps, resulting in an angle θ sampled uniformly between 0 and 2π , that defines the direction of the next flight. Note that due to physical limitations, robots need a finite time to rotate the sampled angle. The number of steps needed to rotate such that the angle of the next flight is equal to the uniformly sampled angle, is computed using the differential steering model [20].

(c) if a robot collides with an object (e.g. another robot), it truncates its motion in both (a) and (b) and instead resorts to the colli-

Algorithm 1: (color online) Overview of the collective Lévy walk. Purple pseudo-code implements the adaptive collective Lévy walk, which is not executed when not considered (e.g. in no patchy environments).

Initialize: Initialize robot swarm (see section 4)

while searching **do**

 Lévy walk (see section 3.2.1) with $\alpha = 1 \rightarrow$ sample ℓ_t

if adaptive collective Lévy walk **then**

if target detected **then**

 | Set $\alpha = 2 \triangleright$ (Brownian motion)

else

 | Decrease α according to eq. 5

end

end

if $\ell_t \geq L_0$ **then** \triangleright long flight

 | Broadcast to local neighborhood

else \triangleright short flight

 Gather information in local neighborhood $n \in \mathcal{S}_t$

if $\exists n \in \mathcal{S}_t \wedge n \notin \mathcal{S}_{t-1}$ **then** \triangleright potential field

 | Rotate with repulsive angle θ_{rep} \triangleright Eq. (4)

end

end

end

sion avoidance procedure. In that procedure, the robot avoids obstacles based on its sensory information, resulting in either a soft turn (slowly moving while turning) or a hard turn (rotate to the opposite directions). The decision between soft and hard turns is based on the free space available in the robot's surroundings. Collision avoidance ensures that each robot is at least a distance d_b from each other robot within the swarm. Once the collision has been avoided, the robot re-samples a flight length and an angle and follows the standard Lévy walk as in (a) and (b).

3.2.2 Collective Lévy walk

As we shall see in Section 4, the ILW does not result in the desired characteristics of the swarm exploration pattern. Specifically, the swarm does not display flight length distributions described by a power-law. In particular, the truncation of long flights suppresses the heavy-tailedness of the collective distribution, while the truncation of short flights does not. Therefore, building upon previous work [17], the purpose of the CLW is to steer robots that are currently not in a long flight away from others which are, so that the truncation probability of long flights due to physical collision decreases. We can ensure this by enabling local communication between individual robots. Local communication is defined through a communication radius r_c , which is larger than the detection radius r_t of each robot.

For each individual, Lévy walks are executed as described above. We define long flights as flight lengths larger than a threshold $\ell \geq L_0$. Then, individuals which are currently following a long flight trajectory broadcast their state, i.e. signaling they are performing a long flight, while additionally ignoring all incoming messages from potential others. Their signal is exploited by other robots within their communication range r_c to derive the relative position of the signal broadcasting robot $\vec{x} = (x_1, x_2)$. Note that this is a relative position, i.e. robots do not have access to any global system for localization, hence preserving the definition of robot swarms. In contrast, individuals which do not follow a long flight gather all incoming messages at time t and compute a repulsive force \vec{F}_{rep} originating

from an artificial potential field [28] as

$$\vec{F}_{rep} = \frac{1}{n} \sum_{j=1}^n \vec{F}^{(j)} \quad (2)$$

$$\text{where } F_i^{(j)} = \frac{-x_i^{(j)}}{d_j^3} \left(\frac{1}{d_j} - \frac{1}{r_c} \right), \quad (3)$$

with n the number of neighbors, $x_i^{(j)}$ the i th component of j th neighboring robot, $i = 1, 2$ the dimensions of the system, and $d_j = \|\vec{x} - \vec{x}^{(j)}\|$ the distance between the robot and the j th neighboring robot which is executing a long flight, i.e. $\ell_j \geq L_0$. However, since the motion of the robots is defined by moving fixed step sizes (i.e. constant velocity) at each time increment, they cannot be steered by exerting this repulsive force onto them. A full continuous description, that evolves robots according to a Langevin equation is considered future work. However, in cases of discretized time, we simply rotate the robot with the angle of the repulsive force

$$\theta_{rep} = \sum_{j=1}^n \text{atan2} \left(\frac{x_2^{(j)}}{x_1^{(j)}} \right), \quad (4)$$

where $\text{atan2}(y/x)$ denotes the 2-argument arc tangent, which properly returns the angle of a vector in all four polar quadrants. The angle of rotation will only be recomputed at $t + 1$ if the set of neighbors contains a new neighbor not present at t . Formally, this means that for a set \mathcal{S}_{t-1} that contains all neighbors at time $t - 1$, we have that the repulsive angle will only be recomputed if there exists a neighbor $n \in \mathcal{S}_t \wedge n \notin \mathcal{S}_{t-1}$. As soon as the robot with the long flight moves out of the communication range of these robots, they resort again to their own Lévy walk procedure described above.

3.2.3 Adaptive collective Lévy walk

As mentioned above, the patchy target distribution agrees with a large number of realistic environments. In such environments, the patches are usually much sparser than the targets within a specific patch. Therefore, in most works, the targets of a particular patch are modeled with a high density and are distributed uniformly. For such patchy target distributions, we propose an adaptive exploration mechanism that extends the collective Lévy walk and improves the swarm search efficiency. This is achieved by adapting the Lévy parameter α based on a local density estimation that is performed by the individual robot. Namely, whenever a target is detected (hence a patch is discovered), each robot can assume that another target is nearby, i.e. the density of targets within the patch is relatively high. Since Brownian motion is optimal in environments with dense target distributions [2, 35], we propose to change the parameter such that $\alpha = 2$ if targets are detected, and decreases to $\alpha = 1$ if no targets have been detected for a specific time period. To ensure this, the Lévy parameter α depends on a discrete time counter τ that increments by 1 each simulation tick wherein a robot has not detected a target. Whenever a target is detected, the counter is reset ($\tau = 0$). We define the Lévy parameter of each robot as

$$\alpha = \max(1, \text{erfc}[\beta(\tau - C)]), \quad (5)$$

where $\text{erfc}(x)$ is the complementary error function and $C = 50$ an empirically chosen constant that determines slope (see figure 1a), that represents the time threshold within which the robot's motion transfers from Brownian back to Lévy. Furthermore, we set $\beta = 0.04$. When $\tau = 0$, we have $\alpha = \text{erfc}(-2) = 2$, and thus

we execute Brownian motion. As τ increases, the walk parameter decreases to $\alpha = 1$ at $\tau = C$. For cases where $\tau > C$, the complementary error function is smaller than 1, hence we ensure $1 \leq \alpha \leq 2$ by the max function. Note that adapting α according to equation 5 can still be applied to environments wherein targets are not heterogeneously distributed, since each individual will return to execute Lévy walks when far from any target.

3.3 Collective measures

After discussing the individual behavior of each robot, let us now discuss the global metrics emerging from the individual behaviors. Our aim is to analyze both the collective trajectory in order to examine whether the individual exploration behavior translates towards similar collective behavior, as well as the actual search efficiency of the robot swarm.

3.3.1 Collective trajectory

We measure the distribution over flight lengths of the swarm as a whole, in order to verify if it still obeys a power-law distribution. Hence, we gather instances of flight lengths, which include both the ones that were executed without truncation and those which were truncated due to collisions. Since we gather flight lengths for each individual and pool them to compute a single distribution, we call this the *collective distribution*. The analysis of this distribution is executed according to Clauset *et al.* [8], which is implemented in the Python package `powerlaw` [1].

The procedure is summarized as follows. First, we fit a power-law to the collective distribution and measure the goodness of fit, i.e. how well a power-law fits the data in comparison to a different distribution. Since distributions are by definition heavy-tailed if the tail is not exponentially bounded, we conclude the distribution is not a power-law if the exponential distribution is a better fit. To identify the power-law as a better fit, we compute the log-likelihood ratio (LLR) between a power-law and an exponential distribution. Hence, if the sign of the LLR is positive, a power-law better fits the collective distribution and vice versa if the LLR is zero or negative. Additionally, we compute a confidence p -value with a significance level of 0.05 [1], ascribing a confidence to the LLR. We compare the individual Lévy walk with the collective Lévy walk for various swarm sizes N .

3.3.2 Collective search efficiency

The distribution over flight lengths and the target distribution both affect the efficiency of searching for targets. Congruent with previous works on individual foragers, we define the search efficiency as the total number of targets k found by the swarm within a certain traveled distance. However, while previous works considered only a single forager [31, 3], here we adopt a search efficiency more appropriate for swarms. Let us define the swarm search efficiency η as the average search efficiency of each individual:

$$\eta = \frac{1}{N} \sum_{i=1}^N \frac{m_i}{\mathcal{L}_i}, \quad (6)$$

where \mathcal{L}_i is the total distance traveled, and m_i the number of targets detected, by each individual i within a swarm of size N . This definition of the search efficiency indicates the frequency of target detection by the swarm, i.e. how much distance needs to be traveled in order to find a target. Since rotation generally takes less energy than




Figure 1: (color online) (a) The adaptive Lévy parameter α in the ACLW, as a function of the number of steps wherein no target is detected τ . Specific parameters are $C = 50$ and $\beta = 0.04$. (b) Screenshot of the ARGoS setup used in the experiments. The blue robots are the footbots, and the green circles represent the targets that can be detected by the ground sensors of each robot.

walking, the traversed distance is the main indicator of efficiency for the random search [31, 21]. The traveled distance of the swarm differs depending on the individual behavior, therefore all three distinct methods carry different search efficiencies.

Additionally, we measure the fraction of targets detected by the swarm, i.e. the number of *distinct* (unique) targets u found by the swarm divided by the total number of targets distributed over the environment,

$$\epsilon = \frac{u}{M}. \quad (7)$$

This allows us to study the dispersivity of the robot swarm, indicating how well the swarm distributes its individuals over the environment to detect as many unique targets as possible.

Lastly, in heterogeneous environments, we define the patch search efficiency, which indicates how many targets have been detected within each patch. We extend the definition given by Wosniack *et al.* [36] to accommodate for multiple foragers. Hence, we define the patch search efficiency ν as

$$\nu = \frac{1}{\mathcal{L}} \frac{k_p}{K_p} \sum_{k=1}^{k_p} \frac{m_k}{1 + \frac{|\bar{m} - m_k|}{m_k}}, \quad (8)$$

where $\mathcal{L} = \sum_{i=1}^N \mathcal{L}_i$ the total distance traveled by the swarm, k_p the total number of distinct patches visited by the swarm, K_p the total number of patches, m_k the number of targets detected in patch k and \bar{m} the average number of targets detected per patch. Note that the homogeneous environment can be recovered by having $M_i = 1$, effectively placing only a single target within each patch. In that case, the patch search efficiency equals the search efficiency η . The denominator of ν is minimized if all patches are equally explored, i.e. if $m_k = \bar{m}$. While we do not directly study different patch qualities, the patch search efficiency ν effectively measures the balancedness of the target detection, where optimization of ν directly relates to visiting a diverse set of patches.

4 Results and discussion

We simulate the individual behaviors described in the previous section in an environment with periodic boundary conditions. Simulations are executed in the multi-robot, physics-based simulator ARGoS [22] (see figure 1b). We increment simulations with steps $\Delta t = 0.1$ s, and run experiments for $T = 5 \cdot 10^4$ simulation




Figure 2: Analysis of the distribution over flight lengths, for various swarm sizes N . Error bars represent one standard deviation. (a) The log-likelihood ratio (LLR) between a heavy-tailed and an exponential distribution. (b) Corresponding p -value of the LLR, dotted line corresponds to $p = 0.05$. (c) The estimated Lévy parameter α that best describes the fitted power-law.

ticks (1 s = 10 simulation ticks). All experiments are executed using simulated footbot robots that have a target detection radius $r_t = d_b = 17$ cm. Each footbot robot is equipped with 24 proximity sensors for obstacle avoidance, range-and-bearing actuators to send messages of at most 10 bytes, range-and-bearing sensors to receive messages and ground sensors for target detection. Specifically, the ground sensor detects the floor color, where targets are designed such that the robot receives a 1 input when it is positioned over the target. Note that communication using the range-and-bearing sensors is only possible when the footbots are in each others line-of-sight. The size of the robot’s communication range, within which information is broadcast and received is set to $r_c = 135$ cm. For the random walk, we define a long flight threshold L_0 by the number of simulation ticks $T_L = 80$, where flights that take more than T_L simulation ticks are considered long flights. Each robot moves with a fixed velocity $v = 0.17$ m/s, equal to a single search radius per second. All experiments are averaged over 30 seeds, unless mentioned otherwise.

Similar to real-life search tasks, robots are initially placed at a deployment location (i.e. a nest), i.e. initializing the swarm by tightly packing the robots in the center of the arena. As mentioned in Section 3.3, the findings of our experiments capture two collective measures. The first measure investigates the diffusion pattern of the emergent collective trajectory, whereas the second measure characterizes the collective performance of the swarm in terms of its search efficiency.

4.1 Collective flight length distribution

First, we discuss the distribution over flight lengths in the two distinct walks, the ILW and the CLW. Since the main difference in the underlying individual behavior between the CLW and the ACLW is the Lévy parameter (i.e. fixed vs. adaptive), we only discuss the ACLW in the context of a random search. For various swarm sizes N , we measure the LLR and its corresponding p -value (see section 3.3.1) to indicate how well the distribution of the flight length is described by a power-law. We select a significance level of the p -value at $p = 0.05$. Random walks are analyzed in an $L \times L$ environment, where $L = 20$ m. This is chosen to accommodate relatively dense swarms, since low densities can result in heavy-tailed distributions for flight lengths even without applying a specific collective Lévy walk due to absence of collisions. Nevertheless, note that our results do generalize towards large environments, however in such cases one needs to study huge swarm sizes ($N \rightarrow 10^4$), which requires much longer computation times. Hence, we have chosen to demonstrate the

effectiveness of the CLW in smaller environments in order to move computation times within feasible domains.

Results are shown in figure 2, and indicate that the CLW indeed results in heavy-tailed distributions up to larger swarm sizes in comparison to the ILW. This is based on the ratio test that assesses the goodness of fit of the two competing models: a power-law (see eq. 1) and an exponential distribution. When the LLR carries a positive sign, and the coinciding p -value falls below the chosen significance level ($p < 0.05$, as indicated by the dotted line in figure 2b), it implies that a heavy-tailed distribution is more likely to describe the data than the exponential distribution. Indeed, our results indicate that larger swarm sizes, i.e. higher swarm densities, can still result in heavy-tailed distributions for the flight lengths when using the CLW instead of the ILW.

Specifically, for the ILW the p -value starts to exceed the chosen significance level from about $N = 500$ individuals. Whereas, the CLW preserves the heavy-tailed distribution of the flight lengths up to the swarm size of approximately $N = 700$ individuals (see figures 2a and 2b). Consequently, we can state that the CLW better maintains a heavy-tailed collective trajectory for larger swarm sizes. This reflects the desired increase in scalability to execute Lévy walks in environments that contain a higher number of individuals. Such scalability enhances the swarm exploration in large and unknown environments through enabling more diffusive behavior and thus aid the collective foraging process (see section 4.2).

Additionally, we have investigated the Lévy parameter α that best fits the collective distribution over flight lengths in figure 2c. This figure shows that collisions, which occur due to high swarm densities, result in more Brownian-like motion, contrasting with the more optimal Lévy walks. This happens even though the underlying step length sampling procedure used $\alpha = 1$. As an effect, the dispersion characteristics of the swarm change to less diffusive behavior, hindering potential Lévy walks by each individual. Furthermore, our results show that Brownian motion ($\alpha \geq 2$) is achieved for much smaller swarm sizes when the individual behavior is defined by the ILW. Thus, the individual communication captured in the CLW enables larger swarm sizes to display Lévy walks ($\alpha < 2$). Hence, the CLW expands the range of swarm sizes, over which a heavy-tailed distribution fits the distribution over actual (truncated) step lengths.

It is worth noting that the fitted Lévy parameter α does not necessarily reflect the actual value of the sampling procedure. Indeed, following the generalized central limit theorem, the actual flight lengths




Figure 3: (color online) Search efficiency η for different swarm sizes N , in homogeneous and heterogeneous environments. Error bars represent one standard deviation.

do not follow a power-law with $\alpha > 2$, but rather appear to be sampled from a normal distribution. Hence, values above the critical value $\alpha^* = 2$ indicate that flight lengths will not be sampled from a power-law, as reflected in the LLR (see figure 2a). It can be seen that the transition from positive ratios to zero or negative ratios indeed coincides with swarm sizes for which holds that the best fitted approximation of the Lévy parameter $\alpha > 2$. Moreover, since the p -value increases over the swarm size (see figure 2b), the ratio test performed in this study is not enough to draw definite conclusions of the distribution nature. In other words, whether the distribution of the flight lengths is heavy-tailed, exponential or a different distribution for which tests have not been computed remains unconfirmed.

4.2 Collective search efficiency

Next, we study all three distinct individual behaviors, the ILW, CLW and ACLW, in both homo- and heterogeneous environments (see section 3.1). Here, to ensure target sparsity, we set the environment size to $L = 1000r_t$. However, as previously noted in section 3.3.1, when the environment size increases, the effectiveness of the proposed methods CLW and ACLW decreases. Henceforth, we initialize the swarm in a more realistic setting, which tightly packs the individuals in the center of the environment, e.g. a nest. This implies that the initial swarm density is high around the environment center, and hence the CLW (and thus the ACLW) methods are in place to increase the dispersive behavior of the swarm.

In the homogeneous environments, $M = 100$ point-like targets are uniformly distributed over the environment. Hence, the total target density

$$\rho = \frac{M\pi r_t^2}{L^2} = \pi \cdot 10^{-4} \sim 10^{-4} \text{ m}^{-2},$$

is low such that the target distribution can be considered sparse. In the heterogeneous case, we follow the argument of Nurzaman *et al.* [21], and deploy $K_p = 10$ patches with radius $R = 0.01L = 10r_t$ to ensure target sparsity. However, within the patches, the density should be high, thus within each single patch we distribute $M_p = 100$ targets, resulting in $M = K_p \cdot M_p = 1000$ targets to be found within the patchy environment.

4.2.1 Target search efficiency

Let us discuss both the homo- and the heterogeneous environments and the target search efficiency of the swarm (see section 4.2). We measure the search efficiency η (see Eq. (6)) for varying swarm sizes,




Figure 4: (color online) Fraction of targets detected ϵ for different swarm sizes N , in homogeneous and heterogeneous environments. Error bars represent one standard deviation.

which results in the trends shown in figure 3. The search efficiency increases as the swarm size increases, with both the ACLW and the CLW outperforming the ILW.

First, we discuss homogeneous environments. In these cases, the CLW displays higher search efficiency than the ILW. This is the direct result of the increase in diffusivity, due to the long flights being less rare since other robots move out of the way of the robot currently in a long flight. Note that for small swarm sizes, the performance difference between the three methods is almost negligible. This is unsurprising, since long flight truncation due to collisions are less likely when the swarm is relatively small. Hence, the search efficiency η is comparable in those cases. We further note that ϵ , the fraction of unique targets found, displays a very similar dependency on the swarm size, regardless of which individual behavior is selected. This means that even though the search efficiency is increased, the total number of different targets detected is approximately the same, with the CLW and the ACLW slightly outperforming the ILW (see figure 4a). Thus, the highest contribution to the increase in search efficiency stems from revisiting targets.

Finally, even in homogeneous environments, the ACLW shows higher search efficiencies than the CLW, however note that the fraction of unique targets found is the same. This can be explained, since the proposed adaptation of the Lévy parameter α does not aid the search for targets which are homogeneously distributed, since the assumption that targets are distributed in patches does not hold. However, since α is changed to exhibit Brownian motion, re-visitation of the target is much more likely in the ACLW than in the CLW, resulting in the notable difference in search efficiency (see figure 3a).

In heterogeneous environments, we again notice that the ACLW outperforms the CLW, which in turn outperforms the ILW, but the difference between the ACLW and the CLW and ILW is much more pronounced. Indeed, the efficiency of the swarm to search for targets distributed within the patches is greatly increased when the behavior of each individual becomes adaptive. The result of the change to more Brownian-like motion when targets are detected, is of great benefit to each individual, and hence the search efficiency of the entire swarm increases. First, this is visible in the swarm search efficiency η , which clearly indicates that the ACLW outperforms the CLW due to the fact that once individuals detect a patch, they are much more likely to stay within the patch (i.e. by adapting the Lévy parameter α). Since the patches have high target density, this naturally increases the number of targets found per distance traveled. Second, the fraction of unique targets ϵ increases (see figure 4b), also due to the fact that patches are more thoroughly searched by the individuals that detect a patch.




Figure 5: (color online) Patch search efficiency ν for different swarm sizes N , within the heterogeneous environment with $K_p = 10$ patches each containing $M_p = 100$ targets. Error bars represent one standard deviation.




Figure 6: Number of foragers that have detected a certain amount of targets, for individuals following the adaptive collective Lévy walk (ACLW). Note the breakline, which indicates a change in scale, aiding the visibility of the lower bars. For large swarm sizes, there exist individuals which find a very large fraction of the total amount of targets, while many individuals find a very small amount of targets, or no target at all.

4.2.2 Patch search efficiency

Next, we investigate the patch search efficiency of the proposed methods. Note that the patch search efficiency is only investigated in heterogeneous environments. Results are depicted in figure 5, which indicates a clear increase in efficiency over the ILW by both the CLW and the ACLW. Due to the ACLW being specifically tailored to exploit the fact that targets are distributed within patches, the swarm improves the patch search efficiency by visiting a large number of unique patches. This indicates that, even though in our case each patch carries the same type of target, our adaptive approach provides the swarm with a higher patch diversity. In turn, when a more varied set of detected targets is preferred, the ACLW provides a way to enhance patch search efficiency as well as patch visitation diversity.

We do wish to note that the large variance of both the search efficiency η and the patch search efficiency ν (see figures 3b and 5) does not disappear with increases in the number of different seeds over which the averages are computed. Namely, each realization of the search experiment consists of a few walkers detecting a large amount of targets, whereas most of the swarm detects nothing. Figure 6 shows the number of individuals that found a certain number of targets for the ACLW executing search tasks in the heterogeneous environment. As can be seen in figure 6, the distribution of the number of targets detected is heavily skewed towards the lower range (0-100 targets). Note the presence of a breakline in the figures, introduced to aid the visibility of the smaller bars. As the analysis reveals, in large swarms there exist few individuals which detect a large amount of targets compared to the vast majority detecting very little. This

is a direct result of the collision avoidance and the adaptive motion. Specifically, whenever an individual enters a patch, it is very likely to stay there for a large period of time due to it executing Brownian motion. As a result, the patch cannot be easily visited by other individuals of the swarm due to collision avoidance resulting in repulsion of the particular patch. Hence, the lucky few individuals that happen to locate a patch earlier contribute largely to the (patch) search efficiency, resulting in high variances even when a larger number different seeds are considered. Additionally, this highlights the importance of the adaptation of the Lévy parameters, since individuals that do locate a patch should exploit it as much as possible.

5 Conclusion

To the best of our knowledge, we are the first to have designed an algorithm that improves the search efficiency in foraging tasks by increasing the diffusive capabilities of a robot swarm for higher numbers of individuals. Enabling the swarm to follow heavy-tailed distributions through applying an artificial potential field method originating from local communication, results in a more significant number of large flight lengths being executed by the swarm. The proposed collective Lévy walk outperforms individual sampling of Lévy flight steps which are unreliable in dense swarms due to collision avoidance procedures. Such avoidance behaviors truncate the large flight lengths, which hampers the diffusive capabilities of the collective, illustrated by the step length distribution not being fitted by a heavy-tailed distribution in favour of exponential distributions. This behavior results in more Brownian-like motion, which is undesired in favour of the Lévy walk. For large swarm sizes, the collective walk outperforms naive individual behavior, while carrying similar performance in small swarms.

Furthermore, we have implemented an adaptive procedure specifically tailored to patchy environments. In these environments, our approach enables each individual to adapt its Lévy parameter α according to the assumption that targets are distributed densely within patches. By changing the sampling of flight lengths from Lévy distributions, which are heavy-tailed, to the normal distribution, which is exponentially bounded, each individual is able to remain within the patch for longer. This results in higher search efficiencies in heterogeneous environments, wherein targets are patchily distributed. Nevertheless, the distribution of the detected targets over the swarm individuals is skewed, since few lucky individuals withhold the detected patches for themselves, due to their Brownian motion and obstacle avoidance. Furthermore, the adaptive behavior does not degrade performance in homogeneous environments, indicating that the adaptive behavior is preferred.

In addition to the identified future perspectives within this study, we highlight a few potential follow-up studies. First, we have only studied random walks wherein pure ballistic motion ($\alpha \rightarrow 0$) is omitted. Even though in most cases where sparse target distributions are considered, $\alpha_{opt} = 1$ emerges as the optimal Lévy parameter [31, 3, 34], some definitions of search efficiency are optimized for more ballistic Lévy walks with $0 < \alpha_{opt} < 1$ [36]. Possibly, a mixture of individuals with different Lévy parameter, or a more adaptive structure wherein more ballistic motion is included, can improve search efficiencies over a wider range of environments.

One can also consider heterogeneous environments wherein each patch carries a different relevance, i.e. different qualities for different patches. The quality of patches can furthermore depend on current demand of the swarm, i.e. patches can represent food or potential nesting sites. Optimization of the quality, instead of target detection

efficiency studied in this work, is currently an often overlooked aspect of random searches which deserves more attention.

Additionally, further study of the transition from the collective flights following a power-law to an exponential distribution is of great interest. While recent work has shown that power-laws are rare in general [24, 7], it is of great interest to study cases wherein the swarms can be designed such that the flight lengths exhibit power-law characteristics. In this study, we have provided elementary research into the critical density at which the transition occurs, however further study wherein actual critical components are extracted is necessary for understanding the behavior around the critical density.

The proposed collective Lévy walks are of great benefit for collective systems to search for sparsely distributed targets, and open up future studies into collective Lévy walks for large-scale systems.

ACKNOWLEDGEMENTS

The authors would like to thank Ilja Rausch for the useful discussions and the invaluable help with setting up the experiments.

REFERENCES

- [1] Jeff Alstott, Ed Bullmore, and Dietmar Plenz, ‘powerlaw: a Python package for analysis of heavy-tailed distributions’, *PLoS one*, **9**(1), e85777, (2014).
- [2] F. Bartumeus, J. Catalan, U. L. Fulco, M. L. Lyra, and G. M. Viswanathan, ‘Optimizing the encounter rate in biological interactions: Lévy versus brownian strategies’, *Phys. Rev. Lett.*, **88**, 097901, (2002).
- [3] Frederic Bartumeus, M. G. E. da Luz, G. M. Viswanathan, and J. Catalan, ‘Animal search strategies: a quantitative random-walk analysis’, *Ecology*, **86**(11), 3078–3087, (2005).
- [4] Levent Bayindir, ‘A review of swarm robotics tasks’, *Neurocomputing*, **172**, 292–321, (2016).
- [5] Jacob Beal, ‘Superdiffusive dispersion and mixing of swarms’, *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, **10**(2), 10, (2015).
- [6] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo, ‘Swarm robotics: A review from the swarm engineering perspective’, *Swarm Intelligence*, **7**(1), 1–41, (2013).
- [7] Anna D Broido and Aaron Clauset, ‘Scale-free networks are rare’, *Nature communications*, **10**(1), 1017, (2019).
- [8] Aaron Clauset, Cosma Rohilla Shalizi, and M. E. J. Newman, ‘Power-law distributions in empirical data’, *SIAM Review*, **51**(4), 661–703, (2009).
- [9] Cristina Dimidov, Giuseppe Oriolo, and Vito Trianni, ‘Random walks in swarm robotics: an experiment with kilobots’, in *International Conference on Swarm Intelligence*, pp. 185–196. Springer, (2016).
- [10] Andrea Falcón-Cortés, Denis Boyer, and Gabriel Ramos-Fernández, ‘Collective learning from individual experiences and information transfer during group foraging’, *Journal of the Royal Society Interface*, **16**(151), 20180803, (2019).
- [11] AS Ferreira, EP Raposo, GM Viswanathan, and MGE Da Luz, ‘The influence of the environment on Lévy random search efficiency: fractality and memory effects’, *Physica A: Statistical Mechanics and its Applications*, **391**(11), 3234–3246, (2012).
- [12] Heiko Hamann, *Swarm robotics: A formal approach*, Springer, 2018.
- [13] Sonia Kéfi, Max Rietkerk, Concepción L Alados, Yolanda Pueyo, Vasilios P Papanastasis, Ahmed ElAich, and Peter C De Ruiter, ‘Spatial vegetation patterns and imminent desertification in mediterranean arid ecosystems’, *Nature*, **449**(7159), 213, (2007).
- [14] Miquel Kegeleirs, David Garzón Ramos, and Mauro Birattari, ‘Random walk exploration for swarm mapping’, in *Annual Conference Towards Autonomous Robotic Systems*, pp. 211–222. Springer, (2019).
- [15] Yara Khaluf, ‘Edge detection in static and dynamic environments using robot swarms’, in *2017 IEEE 11th International Conference on Self-Adaptive and Self-Organizing Systems (SASO)*, pp. 81–90. IEEE, (2017).
- [16] Yara Khaluf and Heiko Hamann, ‘Modulating interaction times in an artificial society of robots’, in *The 2018 Conference on Artificial Life: A Hybrid of the European Conference on Artificial Life (ECAL) and the International Conference on the Synthesis and Simulation of Living Systems (ALIFE)*, pp. 372–379. MIT Press, (2019).
- [17] Yara Khaluf, Stef Van Havermaet, and Pieter Simoens, ‘Collective Lévy walk for efficient exploration in unknown environments’, in *Artificial Intelligence: Methodology, Systems, and Applications*, eds., Gennady Agre, Josef van Genabith, and Thierry Declerck, pp. 260–264, Cham, (2018). Springer International Publishing.
- [18] Yara Khaluf, Seppe Vanhee, and Pieter Simoens, ‘Local ant system for allocating robot swarms to time-constrained tasks’, *Journal of Computational Science*, **31**, 33–44, (2019).
- [19] Simon A. Levin, ‘Multiple scales and the maintenance of biodiversity’, *Ecosystems*, **3**(6), 498–506, (Nov 2000).
- [20] GW Lucas, ‘A tutorial and elementary trajectory model for the differential steering system of robot wheel actuators’, *The Rossum Project*, (2001).
- [21] Surya G Nurzaman, Yoshio Matsumoto, Yutaka Nakamura, Kazumichi Shirai, Satoshi Koizumi, and Hiroshi Ishiguro, ‘An adaptive switching behavior between Lévy and Brownian random search in a mobile robot based on biological fluctuation’, in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1927–1934. IEEE, (2010).
- [22] Carlo Pinciroli, Vito Trianni, Rehan OGrady, Giovanni Pini, Arne Brutschy, Manuele Brambilla, Nithin Mathews, Eliseo Ferrante, Gianni Di Caro, Frederick Ducatelle, et al., ‘ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems’, *Swarm intelligence*, **6**(4), 271–295, (2012).
- [23] E. P. Raposo, Sergey V. Buldyrev, M. G. E. da Luz, M. C. Santos, H. Eugene Stanley, and G. M. Viswanathan, ‘Dynamical robustness of lévy search strategies’, *Phys. Rev. Lett.*, **91**, 240601, (2003).
- [24] Ilja Rausch, Yara Khaluf, and Pieter Simoens, ‘Scale-free features in collective robot foraging’, *Applied Sciences*, **9**(13), 2667, (2019).
- [25] MC Santos, EP Raposo, GM Viswanathan, and MGE Da Luz, ‘Optimal random searches of revisitable targets: crossover from superdiffusive to ballistic random walks’, *EPL (Europhysics Letters)*, **67**(5), 734, (2004).
- [26] Onur Soysal and Erol Şahin, ‘A macroscopic model for self-organized aggregation in swarm robotic systems’, in *International Workshop on Swarm Robotics*, pp. 27–42. Springer, (2006).
- [27] D.J.T Sumpter, ‘The principles of collective animal behaviour’, *Philosophical Transactions of the Royal Society B: Biological Sciences*, **361**(1465), 5–22, (2006).
- [28] Donny Sutantyo, Paul Levi, Christoph Möslinger, and Mark Read, ‘Collective-adaptive lévy flight for underwater multi-robot exploration’, in *2013 IEEE International Conference on Mechatronics and Automation*, pp. 456–462. IEEE, (2013).
- [29] V. Trianni, *Evolutionary Swarm Robotics*, Springer-Verlag, 2008.
- [30] Gabriele Valentini, Eliseo Ferrante, Heiko Hamann, and Marco Dorigo, ‘Collective decision with 100 kilobots: Speed versus accuracy in binary discrimination problems’, *Autonomous Agents and Multi-Agent Systems*, **30**(3), 553–580, (2016).
- [31] Gandimohan M Viswanathan, Sergey V Buldyrev, Shlomo Havlin, MGE Da Luz, EP Raposo, and H Eugene Stanley, ‘Optimizing the success of random searches’, *Nature*, **401**(6756), 911, (1999).
- [32] Bang Wang, Hock Beng Lim, and Di Ma, ‘A survey of movement strategies for improving network coverage in wireless sensor networks’, *Computer Communications*, **32**(13-14), 1427–1436, (2009).
- [33] Henri Weimerskirch, ‘Are seabirds foraging for unpredictable resources?’, *Deep Sea Research Part II: Topical Studies in Oceanography*, **54**(3), 211 – 223, (2007).
- [34] M. E. Wosniack, M. C. Santos, E. P. Raposo, G. M. Viswanathan, and M. G. E. da Luz, ‘Robustness of optimal random searches in fragmented environments’, *Phys. Rev. E*, **91**, 052119, (2015).
- [35] Marina E Wosniack, Marcos C Santos, Ernesto P Raposo, Gandhi M Viswanathan, and Marcos GE da Luz, ‘The evolutionary origins of Lévy walk foraging’, *PLoS computational biology*, **13**(10), e1005774, (2017).
- [36] ME Wosniack, EP Raposo, GM Viswanathan, and MGE da Luz, ‘Efficient search of multiple types of targets’, *Physical Review E*, **92**(6), 062135, (2015).
- [37] V Zaburdaev, S Denisov, and J Klafter, ‘Lévy walks’, *Reviews of Modern Physics*, **87**(2), 483, (2015).
- [38] Kun Zhao, Raja Jurdak, Jiajun Liu, David Westcott, Branislav Kusy, Hazel Parry, Philipp Sommer, and Adam McKeown, ‘Optimal Lévy-flight foraging in a finite landscape’, *Journal of the Royal Society Interface*, **12**(104), 20141158, (2015).