

# Towards an Adaptive Lévy Walk Using Artificial Endocrine Systems

Hugo Sardinha<sup>†\*</sup>, Mauro Dragone<sup>\*</sup> and Patricia A. Vargas<sup>†</sup>

<sup>\*</sup>School of Engineering and Physical Sciences

<sup>†</sup>School of Mathematical and Computer Sciences

Edinburgh Centre for Robotics, Heriot-Watt University

Email: {hs20, m.dragone, p.a.vargas}@hw.ac.uk

**Abstract**—Behavioural adaptation is often observed in foraging animals via coupling periods of localized search and long straight forward motions, a well-known strategy named *Lévy Walk*. In this paper we propose an adaptive *Lévy Walk* model to control an autonomous agent. The model is comprised of a Lévy-based controller modulated by an artificial endocrine system optimised through evolutionary techniques. This new approach enables the agent to control the transition between localized search and long relocation when exposed to external stimulus. The model is tested in exploration tasks where environments have resources clustered into patches. Further tests incorporated environments with different patch characteristics, such as patch size or resource distribution within patches. Our model has shown to outperform the benchmark approach in terms of search efficiency, highlighting the benefits of combining a Lévy Walk based controller with a biologically inspired strategy for adaptation.

**Keywords**—*Artificial Endocrine Systems; Adaptation; Lévy Walk; Biologically Inspired Algorithms; Foraging; Autonomous Agents.*

## I. INTRODUCTION

Foraging animals in nature have long been observed to exhibit adaptive search strategies in order to find beneficial environmental conditions, such as water or food sources [1]. To describe this motion of animals in the wild, computational ecologists firstly attempted to use Brownian models to fit empirical data, over large spatial scales and long temporal scales, with relative success [2]. However, these uncorrelated models of motion did not account for the tendency that animals show to continue moving in the same direction [3]. To better describe this behaviour, two major models have been employed, the Correlated Random Walk (CRW) [4] and the Lévy Walk (LW) [5]. As in Brownian motion, the CRW considers a Gaussian distribution of walk lengths, but unlike its predecessor, draws the re-orientation angle from a non-uniform distribution, centered around the current heading, thus generating a directional persistence. Conversely, the LW model considers an uniform distribution to draw new orientations but, unlike the previous models, draws length walks from a power law distribution instead. Since power laws are heavy tailed, this leads to occasional *long* walks, effectively coupling periods of localized random search with periods of ballistic relocation across the domain [6]. Figure 1 shows an example of trajectories generated by Brownian and Lévy motions. Consequently, several works used these models to develop robot controllers in foraging scenarios, corroborating theoretical results [7] [8]. In fact, LWs have become regarded as the optimal foraging strategy by subsequent works [9], particularly when points of interest, or resources, are sparse and randomly distributed [10]. However, different scenarios exist where such resources may not materialize in this manner; for example, in agricultural tasks such as the mapping of weeds in a field, where resources (weeds) tend to be clustered

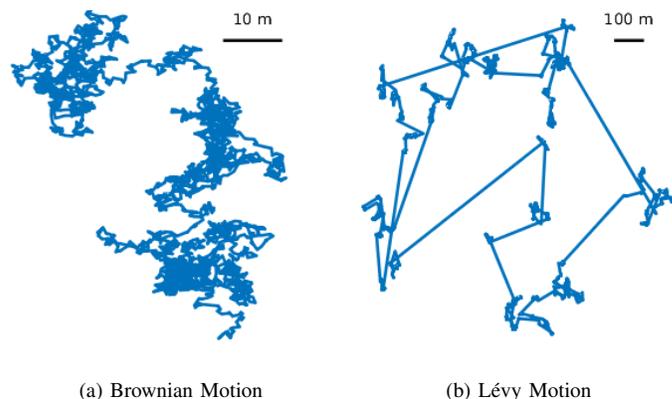


Figure 1. Examples of trajectories generated by Brownian and Lévy motions

together in *patches* [11]. Danchin’s [12] definition of a *patch* —“an homogeneous resource containing area (or part of habitat) separated from others by areas containing little or no resources”—is of crucial importance for the solution we here present. Since resources are not distributed uniformly, but exist in regions of locally high density there is an imposing drive for an autonomous agent to be able to adapt its behaviour in order to find them efficiently [13]. In fact, one could interpret LWs observed in nature as a consequence of an underlying adaption mechanism, switching to localized search when inside a patch, and switching back to ballistic motions once no new resources are found. To interpret adaptation in the context of either natural or artificial autonomous agents we recall Ashby’s definition by which *a form of behaviour is considered adaptive if it maintains the essential variables within physiological limits* [14]. Such an ability to maintain an internal equilibrium, known as *homeostasis*, is ubiquitous in the animal kingdom [15], and depends strongly on the endocrine systems these animals developed over the course of their evolution [16]. These systems are responsible for the production of hormones that regulate a myriad of bodily functions such temperature, heart-rate, or the desire to hunt, and have also inspired researchers to synthesise Artificial Endocrine Systems (AES) for robot control [17]. Drawing inspiration from hormone-based regulation in natural systems, this work proposes an Endocrine-based Lévy Walk (ELW) model for foraging in patchy environments, by which the observable Lévy process is modulated by an underlying AES endowing the forager with the ability to change its behaviour in the presence of favourable environmental conditions. The paper is divided as follows: Section II highlights the most relevant models for our work; Section III describes in detail our proposed model; Section IV presents our results and how they compare to the benchmark, and finally Section V summarizes our findings.

## II. RANDOM WALK MODELS

Random walks are a class of stochastic processes used to model empirical data. Existing processes arise from the intuitive idea of taking successive steps, each in random direction. Therefore, they consist of displacement events (i.e., walk lengths) interspersed by reorientation events [18]. This section explores the basic concepts of Lévy walks and, in particular, the biological fluctuation method proposed by Nurzaman *et al.* [19], which will be used as a benchmark for our model.

### A. Lévy walks

Lévy Walks are a random walk process characterized by drawing each step length ( $l_s$ ) from a power law distribution:

$$P(l_s) \sim l_s^{-\mu}, \quad 1 < \mu \leq 3 \quad (1)$$

where the parameter  $\mu$  controls the shape of the distribution's tail, making ballistic relocations more (or less) common. Previous work has shown that for  $\mu \geq 3$  the motion becomes Brownian, whereas when  $\mu \rightarrow 1$ , it becomes a series of straight motions with negligible local searches [6]. At each reorientation step, a new heading  $\tau$  is generated such that  $\tau \sim \mathcal{U}(-\pi, \pi)$ . Some authors have proposed replacing an uniform turning angle, and instead use a correlated reorientation, which has shown to produce some improvement in search efficiency, particularly in environments where resources are sparsely distributed [3] [20]. In our work, the proposed adaptive model will be tested both with and without correlation, to study its effect on patchy environments. Correlation is achieved by drawing  $\tau$  from a wrapped Cauchy distribution, whose probability density function, given in [20], is as follows:

$$C(\rho, \tau) = \frac{1}{2\pi} \frac{1 - \rho^2}{1 + \rho^2 - 2\rho \cos(\tau)} \quad (2)$$

where the parameter  $\rho$  represents how correlated the direction of consecutive walks is. On the one hand, when  $\rho = 1$ , correlation is complete and therefore the entire motion is a continuous straight line, while on the other hand, when  $\rho = 0$ , reorientations are in fact not correlated and  $\tau$  assumes an uniform distribution.

### B. Biological Fluctuation

An alternative to distribution-based random walk models was proposed by Nurzaman *et al.*, where the transition between local searches and ballistic motions happens based on the concept of *yuragi* or biological fluctuation [19]. This mechanism is one by which certain bacteria are able to alter their gene expression (and therefore their behaviour) in the presence, or absence, of nutrients. A formal description of such behaviour is given by the attractor selection model, represented by the Langevin equation:

$$\dot{\mathbf{x}}(t) = -\nabla U(\mathbf{x}(t))A(t) + \epsilon(t) \quad (3)$$

where  $\mathbf{x}$  and  $-\nabla U(\mathbf{x}(t))$  are respectively the state and dynamics of the attractor model,  $\epsilon(t)$  is a noise term, and  $A(t)$  represents a variable *activity* which indicates how well the current state fits the environment, chosen in Nurzaman's work to be respectively:

$$U(\mathbf{x}(t)) = (\mathbf{x}(t) - h)^2 \quad (4)$$

$$A(t) = R \cdot A(t-1) + f(t) \quad (5)$$

where  $f(t)$  represents the number of resources sensed, and  $R$  is a decaying coefficient with respect the previous value of  $A(t)$ . The way the system changes from continuously straight motions to local search is modeled by a finite state machine with two states: *swimming* or *gliding*, which corresponds to a forward motion, and *tumbling* which corresponds to a reorientation. One can observe, from the depiction of in Figure 2, that the transition from the *gliding* state  $\mathcal{G}$ , to the *tumbling* state  $\mathcal{T}$  depends on a probability  $P(t)$ . On one hand, if  $P(t)$  is small the gliding motion continues and long relocations are expected, whereas on the other hand, for high values  $P(t)$  a *tumbling* step is more likely to occur, immediately followed by another *gliding* step leading to a local search behaviour.

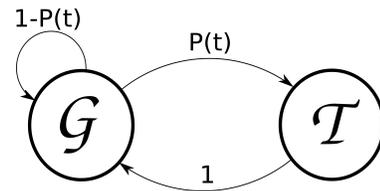


Figure 2. Finite state machine for behaviour

Based on the state of the system, the probability of transitioning between states is computed using (6). One can see, from (3), that when  $A(t)$  increases, the first term becomes dominant and the value of the system's state  $\mathbf{x}$  decreases towards attractor  $h$  (with  $0 < h < 1$ ), therefore leading to a high probability transitioning from  $\mathcal{G}$  to  $\mathcal{T}$ .

$$P(t) = e^{-\mathbf{x}(t)} \quad (6)$$

On the other hand, when  $A(t)$  is very small, the dominant term of (3) is the noise term  $\epsilon(t)$ , which, due to being applied to  $\dot{\mathbf{x}}$ , gradually makes the state  $\mathbf{x}$  diverge from the attractor leading to small value of  $P(t)$ , and therefore a continued gliding motion.

## III. ENDOCRINE-BASED ADAPTIVE LÉVY WALK

The underlying Lévy controller of most (if not all) artificial agents, or foragers, can be considered to have two stages: *Generation* and *Execution*. At the *Generation* stage a tuple  $(\tau, l_s)$  is selected depending on  $\mu$  and, if correlated, on  $\rho$ . This tuple is used in the *Execution* stage to move the agent in the direction  $\tau$  while  $l(t) < l_s$ , where  $l(t)$  is the distance travelled since the beginning of the current walk. Our work proposes that, in order to achieve adaptation, both  $\mu$  and  $l_s$  need to change dynamically according to sensory input. Firstly, as the forager enters a patch it is straightforward to envision that  $\mu$  should increase, so that the behaviour converges to a local search. However, only changing the value of  $\mu$  will have little or no effect if the current step is not completed within the patch in time for another tuple  $(\tau, l_s)$  to be generated. In order to harness this intuition we introduce a *desire* to interrupt the current walk, which will translate to a gradual decrease of  $l_s$  for the ongoing step. The specific AES proposed for Lévy walk adaption is built upon the concepts put forward by Wilson *et al.* [21] and Stradner *et al.* [22], where the level of a hormone  $H$  at time  $t$  can be modelled by the following:

$$H(t) = c_0 + c_1 H(t-1) + c_2 S(t) \quad (7)$$

where  $c_0, c_1, c_2$  are constant coefficients,  $H(t-1)$  represents the previous hormone level and  $S(t)$  is the stimulus received from sensory input. The first term,  $c_0$ , represents a base increment simulating a default and constant hormone production, the second term  $c_1 H(t-1)$  acts as decay over time, and  $c_2 S(t)$  represents the contribution from the sensory stimulus to the overall level of  $H(t)$ . Wilson [21] highlights that one could calculate the settling point of  $H(t)$ , when no stimulus is received, as  $H_s = c_0 / (1 - c_1)$ . Using (7) we model the variation of  $\mu$  as the hormone level itself and define:

$$\mu(t) = a_0 + a_1 \mu(t-1) + a_2 S_\mu(t) \quad (8)$$

where  $S_\mu(t)$  assumes a binary value depending on the variation of number of resources sensed according to (9). Therefore, as the forager enters a patch of resources,  $\mu$  tends to increase, while if there are no new resources  $\mu(t) \rightarrow \mu_s = a_0 / (1 - a_1)$ .

$$S_\mu(t) = \begin{cases} 1, & \Delta f(t) > 0 \\ 0, & \Delta f(t) \leq 0 \end{cases} \quad (9)$$

Modelling the aforementioned *desire* to interrupt the current walk is done in a similar fashion, by considering the hormone level  $\beta$  defined as:

$$\beta(t) = b_1 \beta(t-1) + b_2 S_\beta(t) \quad (10)$$

Note that there is no  $b_0$  term, allowing in fact the value of  $\beta(t)$  to decrease to zero. The stimulus function for  $\beta$  is given in (11), where if  $f(t)$  is increasing, there is no stimulus to the *desire* to interrupt the current walk since this means the current walk step is providing a good strategy to find resources. Conversely, if the  $f(t)$  is decreasing this *desire* increases, and does so proportionally to the normalized value of  $\mu(t)$  between its settling point ( $\mu_s$ ) and its maximum value ( $\bar{\mu} = 3$ ), creating an interdependence of these two artificial hormone quantities as it also the case in several natural systems [16].

$$S_\beta(t) = \begin{cases} 0, & \Delta f(t) \geq 0 \\ \frac{\bar{\mu} - \mu(t)}{\bar{\mu} - \mu_s}, & \Delta f(t) < 0 \end{cases} \quad (11)$$

As one can see, when  $\mu(t) \rightarrow \mu_s$  then the stimulus  $S_\beta \rightarrow 1$  and when  $\mu(t) \rightarrow \bar{\mu}$   $S_\beta \rightarrow 0$ . In practical terms this means that if the forager is finding fewer resources but the its  $\mu(t)$  value is large, it is already performing a local search and the *desire* to interrupt that walk is irrelevant since it would already be a local step. On the other hand, if  $\mu(t)$  is small in the presence of a varying number of resources then the *desire* is relevant and it is stronger as  $\mu(t)$  is further from  $\bar{\mu}$ . Updating the target step length is, in our model simply done by computing (12).

$$l_s = l_s (1 - \beta(t)) \quad (12)$$

In summary, we can consider our endocrine-based model, depicted in Figure 3, to have three main components, namely:

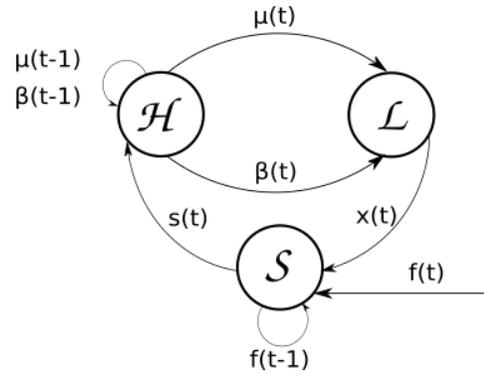


Figure 3. Endocrine-based Lévy Walk model.

the Hormone production module ( $\mathcal{H}$ ) that updates the values of  $\mu(t)$  and  $\beta(t)$ ; the Lévy controller ( $\mathcal{L}$ ) which controls the pose of the forager, generates  $l_s$  depending on  $\mu$  and updates it depending on  $\beta$ ; and finally the sensory input module,  $\mathcal{S}$ , which given the new pose  $x(t)$  and the previously sensed resources updates the stimulus functions.

#### IV. EXPERIMENTS & RESULTS

The performance of our proposed model is assessed by a series of experiments in different environments. The first of which is the one chosen by Nurzaman et al, (Environment I) and comprises an arena of 1000x1000m where 10 patches are randomly distributed. These patches are 10x10m and contain 100 rewards each, also uniformly distributed. As in Nurzaman's work, we consider an agent whose field of view is 2x2m travelling at a speed of 1m/s over 10000s. Furthermore we expand testing to environments where patches are 50x50 (Environment II) and 100x100 (Environment III) maintaining reward density, and use these three environments to test the influence of correlation  $\rho$  in our model. Figure 4 depicts these environments. To choose the parameters for our model, namely the coefficients in (8) and (10), we implemented a Real-Coded Genetic Algorithm (RCGA) where each solution is represented by a generic chromosome of the type:

$$\mu_s \quad a_1 \quad a_2 \quad b_1 \quad b_2 \quad \rho$$

Choosing  $\mu_s$  as a decision variable makes for a more straightforward analysis of the solution, since  $\mu_s$  represents the settling point for  $\mu$ , and  $a_0$  can be obtained from  $\mu_s = a_0 / (1 - a_1)$ . To optimize the model without correlation,  $\rho$  is set to 0 and dropped from the chromosome. Parameters are evolved to maximize search efficiency defined in (13), where  $P$  is the total number of rewards found and  $D$  the total distance travelled. Individuals for crossover are selected by tournament and an arithmetic crossover operator is used.

$$\eta = P/D \quad (13)$$

Mutation is done by selecting a random allele and changing its value randomly between its predefined limits, i.e.,  $\mu_s \in (1, 3]$  and  $a_1, a_2, b_1, b_2, \rho \in [0, 1]$ . The RCGA runs over 250 generations with a randomly initialized population of 100 and a 5% elitism. Results of the optimization process are shown in Table I, with and without correlation. The first observation we can make is that, as the size of patches increases from environment I to III, the value of  $\mu_s$  decreases. In fact, this was an expected result, since in an environment where patches

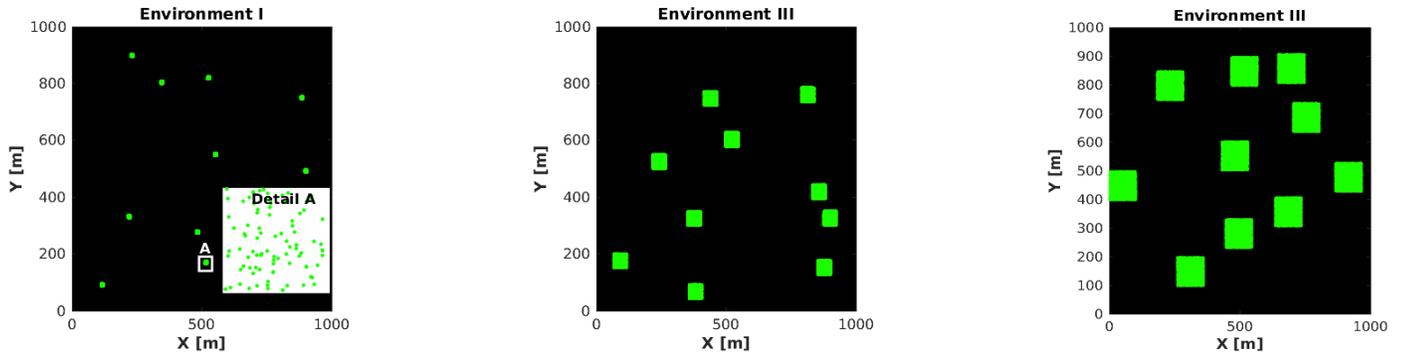


Figure 4. Environments with uniformly distributed rewards.

TABLE I. PARAMETERS OPTIMIZED FOR ENVIRONMENTS I,II AND III.

Environment	$\mu_s$	$a_1$	$a_2$	$b_1$	$b_2$	$\rho$
I	2.344	0.196	0.179	0.995	0.235	—
	2.422	0.216	0.599	0.962	0.424	0.073
II	1.625	0.503	0.254	0.976	0.185	—
	1.809	0.507	0.354	0.917	0.563	0.029
III	1.129	0.841	0.619	0.918	0.191	—
	1.428	0.758	0.681	0.904	0.153	0.077

TABLE II. COMPARISON BETWEEN ELW AND CELW OVER 100 RUNS.

Search Behaviour	Efficiency ( $10^{-2}$ )	Rewards Found ( $10^2$ )	Patches Found
ELW-I	$0.42 \pm 0.12$	$4.18 \pm 1.13$	$7.48 \pm 1.56$
CELW-I	$0.41 \pm 0.15$	$3.98 \pm 1.22$	$7.04 \pm 1.73$
ELW-II	$15.02 \pm 2.14$	$126.50 \pm 11.32$	$9.58 \pm 0.58$
CELW-II	$11.70 \pm 1.70$	$104.24 \pm 13.39$	$9.66 \pm 0.55$
ELW-III	$39.93 \pm 4.78$	$372.39 \pm 42.00$	$9.95 \pm 0.21$
CELW-III	$36.71 \pm 3.48$	$351.60 \pm 31.93$	$9.98 \pm 0.14$

are of considerable size, a strategy that favours more frequent ballistic motions, will tend find new rewards more efficiently. On the other hand, in an environment where patches are very small, more frequent changes of direction are necessary to find these rewards, and therefore the  $\mu_s$  is higher. We also observe that parameter  $a_1$  increases with the increasing size of patches. Recalling (8), the higher  $a_1$  becomes the slower is the decay of  $\mu(t)$ . An increasing value  $a_1$  with patch size shows that it is beneficial to not let the value of  $\mu(t)$  decay too abruptly, sustaining more localized search for a longer period. Concurrently  $a_2$ , the weight of the stimulus to the value of  $\mu(t)$ , also increases. Since  $\mu_s$  is smaller for such environments, the change of  $\mu(t)$  to a point that translates into local search needs to occur at a higher rate, and thus the stimulus has a bigger weight. As for the coefficients that modulate  $\beta$ , we see that  $b_1$  always maintains a high value, but decreases only slightly with increasing patch size, showing that a slow decay of the *desire* to interrupt the current walk is always desirable regardless of the size of patches. As for the stimulus to this *desire*, i.e., the  $b_2$  parameter, has a higher value when patches are smaller. Naturally, given the smaller size of patches, the weight of the stimulus must be stronger so that the current walk is interrupted sooner. Perhaps the more interesting result is the one concerning the correlation  $\rho$ . These results show that  $\rho$  always converges to very small values, hinting that, in patchy environments, directional correlation between steps might not play a significant role. To confirm this hypothesis we show in Table II the average values for *Efficiency*, number of *Rewards Found* and number of *Patches Found*, for both the Endocrine-based Lévy Walk (ELW) and its correlated version (CELW). These results show that, in fact, even a negligible correlation can have an apparent negative impact on the overall performance of the system, since the ELW tends to outperform its correlated counterpart for most

TABLE III. P-VALUE FOR THE 2-SAMPLE KOLMOGOROV-SMIRNOV TEST, FOR ALL METRICS BETWEEN ELW AND CELW BEHAVIOURS.

Search Behaviours	Efficiency	Rewards Found	Patches Found
ELW-I / CELW-I	0.443	0.432	0.893
ELW-II / CELW-II	$\ll 0.01$	$\ll 0.01$	0.987
ELW-III / CELW-III	$\ll 0.01$	$\ll 0.01$	0.998

metrics in all environments. To further investigate the statistical difference between models, we calculate the  $p$ -value for the 2-sample Kolmogorov-Smirnov (KS) goodness-of-fit hypothesis test. For the 2-sample KS test, the null hypothesis is that “two data vectors belong to the same continuous distribution” and therefore, the hypothesis is rejected, at the 99% confidence level, if the  $p$ -value  $< 0.01$ . Table III summarizes the  $p$ -values between each uncorrelated and respective correlated model, for each different metric. These results show that for the models optimized for Environment I, there is no statistical difference in results for any of the metrics, confirming that including correlation, for small patch environments, does not have any impact on the system’s performance. However, for Environments II and III the respective optimized models are in fact statistically different for both the *Efficiency* and *Rewards Found*. In these particular cases it shows that including correlation actually has a detrimental effect to the system’s performance, since metrics’ values shown in Table II are higher for the ELW. Having established that correlation either does not affect or worsens the performance of the system, the next set of results will only compare the ELW and the Yuragi approaches. To also test the flexibility of both approaches, to scenarios where rewards within patches are distributed differently we further extend our analysis to Environments IV and V (Figure 5) where rewards within patches have Gaussian distributions.

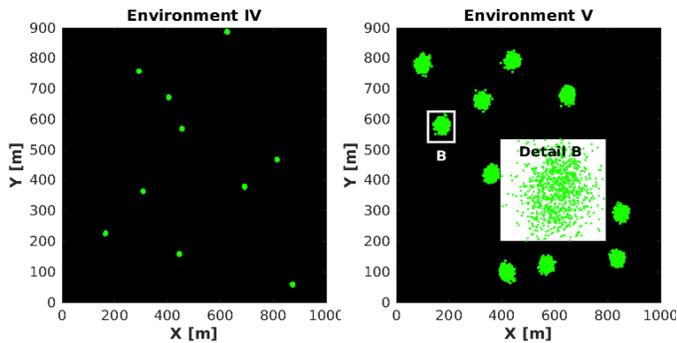


Figure 5. Environments with Gaussian distributed rewards.

TABLE IV. PARAMETERS OPTIMIZED FOR ENVIRONMENTS IV AND V

Environment	$\mu_s$	$a_1$	$a_2$	$b_1$	$b_2$
IV	1.659	0.152	0.511	0.940	0.499
V	1.445	0.489	0.495	0.980	0.194

The parameters of our model for these environments were also obtained via the aforementioned RCGA, and are shown in Table IV. For a fair comparison with the Yuragi model, we conduct the same sensitivity analysis as Nurzaman *et al.* to select the value of  $R$  in (5), which maximizes *efficiency*. These results are summarized in Table V, where we include the values reported by Nurzaman, in brackets, to validate our own implementation. In the following analysis, we refer to Yuragi A/B or C depending on which  $R$  is best suited for a particular environment. Table V shows that Yuragi-A is best suited for Environments II, IV and IV, and Yuragi-B for Environments I and III. We also note that the values obtained with our Yuragi implementation yield very close results to those reported by Nurzaman, validating our implementation for the subsequent analysis. Finally, the comparison between the ELW and Yuragi approaches is presented in Table VI. This table compares the results obtained with the optimized ELW model for each particular environment (ELW-I, ELW-II, etc.) with the corresponding Yuragi model chosen from the sensitivity analysis in Table V. Highlighted values for each metric show that the ELW always yields best performance both across the different metrics and different environments. However, one could still argue that our model requires prior knowledge in order to select optimal parameters for the task. To address this question, we perform a cross-testing analysis and run each ELW model (ELW-I, ELW-II, etc.) in every environment, and compare those results, both with the optimal results for such environment, and with those achieved with the Yuragi approach. This comparison is made in Table VII, where the optimal efficiency values for each environment are highlighted in green and the best Yuragi approach for each environment is highlighted in red. Furthermore, we also highlight the performance obtained with ELW-II and ELW-IV, since these are able to consistently outperform the best Yuragi solution, even in those environments for which the ELW was not specifically optimized. The existence of sets of ELW parameters that lead to a higher performance, in comparison to the Yuragi approach, and regardless of the environment is an important evidence of the superiority of the

TABLE V. SENSITIVITY ANALYSIS ON EFFICIENCY WITH VARYING  $R$ 

Model	$R$	Env I	Env II	Env III	Env IV	Env V
A	0.99	0.26 $\pm$ 0.11 (0.23 $\pm$ 0.11)	<b>8.98<math>\pm</math>2.38</b>	19.74 $\pm$ 3.51	<b>0.35<math>\pm</math>0.13</b>	<b>3.93<math>\pm</math>0.99</b>
B	0.90	<b>0.29<math>\pm</math>0.10</b> (0.28 $\pm$ 0.11)	7.41 $\pm$ 1.98	<b>24.07<math>\pm</math>4.48</b>	0.29 $\pm$ 0.11	3.19 $\pm$ 0.89
C	0.50	0.21 $\pm$ 0.09 (0.17 $\pm$ 0.06)	6.72 $\pm$ 1.61	20.70 $\pm$ 4.69	0.24 $\pm$ 0.10	2.49 $\pm$ 0.79

TABLE VI. METRIC COMPARISON BETWEEN YURAGI AND ELW

Search behaviour	Efficiency ( $10^{-2}$ )	Rewards Found ( $10^2$ )	Patches Found
ELW-I	<b>0.42<math>\pm</math>0.12</b>	<b>4.18<math>\pm</math>1.13</b>	<b>7.48<math>\pm</math>1.56</b>
Yuragi-B	0.29 $\pm$ 0.10	2.61 $\pm$ 1.01	5.20 $\pm$ 1.53
ELW-II	<b>15.02<math>\pm</math>2.14</b>	<b>126.50<math>\pm</math>11.32</b>	<b>9.58<math>\pm</math>0.58</b>
Yuragi-A	8.98 $\pm$ 2.38	64.40 $\pm$ 14.50	6.03 $\pm$ 1.30
ELW-III	<b>39.93<math>\pm</math>4.78</b>	<b>372.39<math>\pm</math>42.00</b>	<b>9.95<math>\pm</math>0.21</b>
Yuragi-B	24.07 $\pm$ 4.48	180.59 $\pm$ 29.03	6.07 $\pm$ 1.48
ELW IV	<b>0.39<math>\pm</math>0.01</b>	<b>3.83<math>\pm</math>0.97</b>	<b>6.89<math>\pm</math>0.97</b>
Yuragi-A	0.35 $\pm$ 0.13	3.03 $\pm$ 1.12	4.45 $\pm$ 1.37
ELW V	<b>6.22<math>\pm</math>0.85</b>	<b>56.71<math>\pm</math>7.13</b>	<b>9.73<math>\pm</math>0.48</b>
Yuragi-A	3.93 $\pm$ 0.99	29.91 $\pm$ 6.67	6.19 $\pm$ 1.41

ELW model, and its inherent increased adaptation ability. It is also interesting to point out that the performance, with respect to every metric, of ELW-II and ELW-IV in environments I,III and V is considerably closer to the optimal value, than to the Yuragi approach, often lying within the standard deviation of the former. This is yet another evidence that improvement in performance can be achieved mostly due to the model itself and that RCGA optimization is mostly useful in fine-tuning.

## V. CONCLUSION AND FUTURE WORK

This work proposes an Endocrine-based Lévy Walk (ELW) model for autonomous agents, that modulates the parameters of a Lévy controller for foraging in patchy environments. Firstly, we analysed the hypothesis that angular correlation between steps can benefit the search efficiency, and have shown that in fact, it either does not have any statistical impact on results, or slightly decreases the efficiency of the model in environments with larger patches. Secondly, we compared our model with a one based on biological fluctuation, across environments where rewards were either uniformly or normally distributed within patches. Both models were compared using three different metrics: *Efficiency*, *Rewards Found* and *Patches Found*. Results show that ELW outperforms the benchmark model in all scenarios. Future work will look into the gap in performance by studying the difference in trajectories, as well as a heatmap analysis of the probability density of visited areas within the environment, and the temporal evolution of hormone levels  $\mu(t)$  and  $\beta(t)$ . We further aim to conduct experiments using real and simulated mobile robots to test how suitable these controllers are to swarm robotics and the control of collective behaviours [23]. We envision that these behaviours could contribute to tackling real-world applications, both in agricultural and environmental domains, such as weed mapping or forest fire prevention, to name a few.

TABLE VII. CROSS-TESTING OF ELW PARAMETERS ACROSS ENVIRONMENTS.

Search Behaviour	Metric	Environment I	Environment II	Environment III	Environment IV	Environment V
ELW-I	Efficiency ( $10^{-2}$ )	0.42±0.12	10.09±1.74	20.64±2.63	0.37±0.15	4.40±0.81
	Rewards Found ( $10^2$ )	4.18±1.13	72.70±9.97	124.68±12.17	3.59±1.30	34.00±5.42
	Patches Found	7.48±1.56	6.01±1.04	3.91±1.73	5.18±1.62	6.83±1.17
ELW-II	Efficiency ( $10^{-2}$ )	0.40±0.10	15.02±2.14	38.81±3.74	0.38±0.11	6.17±0.92
	Rewards Found ( $10^2$ )	3.92±0.99	126.50±11.32	266.51±18.25	3.73±1.05	55.92±7.74
	Patches Found	7.27±1.43	9.58±0.21	8.59±1.08	7.08±1.38	9.80±0.47
ELW-III	Efficiency ( $10^{-2}$ )	0.24±0.07	9.98±1.73	39.93±4.78	0.21±0.07	3.82±0.64
	Rewards Found ( $10^2$ )	2.36±0.69	98.12±16.84	372.39±42.00	2.18±0.78	37.82±6.35
	Patches Found	7.56±1.32	9.85±1.03	9.95±0.21	7.10±1.57	9.93±0.32
ELW-IV	Efficiency ( $10^{-2}$ )	0.41±0.10	12.92±2.02	38.22±3.46	0.39±0.01	5.20±0.79
	Rewards Found ( $10^2$ )	4.16±0.95	113.35±15.87	276.95±18.67	3.83±0.97	48.12±6.85
	Patches Found	7.32±1.29	9.87±0.37	9.49±0.67	6.89±1.29	9.85±0.36
ELW-V	Efficiency ( $10^{-2}$ )	0.36±0.11	15.01±1.87	39.01±3.61	0.35±0.11	6.21±0.85
	Rewards Found ( $10^2$ )	3.60±1.12	127.10±13.51	274.55±17.48	3.56±1.11	56.71±7.13
	Patches Found	6.98±1.50	9.66±0.55	8.87±0.87	7.12±1.45	9.73±0.46
Yuragi-A	Efficiency ( $10^{-2}$ )	0.26±0.11	8.98±2.38	19.74±3.51	0.35±0.13	3.93±0.99
	Rewards Found ( $10^2$ )	2.01±1.01	64.40±14.51	121.61±16.64	3.03±1.12	29.91±6.67
	Patches Found	4.81±1.16	6.03±1.31	3.86±1.4	4.45±1.37	6.19±1.41
Yuragi-B	Efficiency ( $10^{-2}$ )	0.29±0.10	7.41±1.98	24.07±4.48	0.29±0.11	3.19±0.89
	Rewards Found ( $10^2$ )	2.61±1.01	61.80±15.68	180.59±29.03	2.58±0.95	27.17±7.29
	Patches Found	5.20±1.53	7.19±1.41	6.07±1.48	4.85±1.45	7.43±1.65

## REFERENCES

- [1] A. M. Reynolds and C. J. Rhodes, "The lévy flight paradigm: random search patterns and mechanisms," *Ecology*, vol. 90, no. 4, 2009, pp. 877–887.
- [2] H. C. Berg, *Random walks in biology*. Princeton University Press, 93.
- [3] F. Bartumeus, M. G. E. da Luz, G. M. Viswanathan, and J. Catalan, "Animal search strategies: a quantitative random-walk analysis," *Ecology*, vol. 86, no. 11, 2005, pp. 3078–3087.
- [4] P. Bovet and S. Benhamou, "Spatial analysis of animals' movements using a correlated random walk model," *Journal of theoretical biology*, vol. 131, no. 4, 1988, pp. 419–433.
- [5] G. Viswanathan et al., "Lévy flights in random searches," *Physica A: Statistical Mechanics and its Applications*, vol. 282, no. 1-2, 2000, pp. 1–12.
- [6] O. Bénichou, C. Loverdo, M. Moreau, and R. Voituriez, "Intermittent search strategies," *Reviews of Modern Physics*, vol. 83, no. 1, 2011, pp. 82–127.
- [7] G. M. Fricke, F. Asperti-Boursin, J. Hecker, J. Cannon, and M. Moses, "From microbiology to microcontrollers: Robot search patterns inspired by t cell movement," in *Artificial Life Conference Proceedings 13*. MIT Press, 2013, pp. 1009–1016.
- [8] V. Fioriti, F. Fratichini, S. Chiesa, and C. Moriconi, "Levy foraging in a dynamic environment—extending the levy search," *International Journal of Advanced Robotic Systems*, vol. 12, no. 7, 2015, pp. 98–110.
- [9] M. E. Wosniack, M. C. Santos, E. P. Raposo, G. M. Viswanathan, and M. G. da Luz, "The evolutionary origins of lévy walk foraging," *PLoS computational biology*, vol. 13, no. 10, 2017, pp. 1–31.
- [10] V. K. Jandhyala and S. B. Fotopoulos, "Applications of random search methods to foraging in ecological environments and other natural phenomena—a review," *Environmetrics*, vol. 29, no. 5-6, 2018, pp. 1–12.
- [11] F. Castaldi, F. Pelosi, S. Pascucci, and R. Casa, "Assessing the potential of images from unmanned aerial vehicles (uav) to support herbicide patch spraying in maize," *Precision Agriculture*, vol. 18, no. 1, 2017, pp. 76–94.
- [12] E. Danchin, L.-A. Giraldeau, and F. Cézilly, *Behavioural ecology*, 2008.
- [13] J. Wawerla and R. T. Vaughan, "Robot task switching under diminishing returns," in *IEEE/RSJ ICIRS*, 2009, pp. 5033–5038.
- [14] W. Ashby, *Design for a brain: The origin of adaptive behaviour*. Springer Science & Business Media, 2013.
- [15] P. L. deFur, "Use and role of invertebrate models in endocrine disruptor research and testing," *ILAR journal*, vol. 45, no. 4, 2004, pp. 484–493.
- [16] P. A. Vargas, R. C. Muioli, F. J. Von Zuben, and P. Husbands, "Homeostasis and evolution together dealing with novelties and managing disruptions," *Int. J. Intel. Comp. and Cyb.*, 2009, pp. 435–454.
- [17] P. Vargas et al., "Artificial homeostatic system: a novel approach," in *ECAL*. Springer, 2005, pp. 754–764.
- [18] F. Bartumeus, "Lévy processes in animal movement: an evolutionary hypothesis," *Fractals*, vol. 15, no. 02, 2007, pp. 151–162.
- [19] S. G. Nurzaman et al., "An adaptive switching behavior between levy and brownian random search in a mobile robot based on biological fluctuation," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 1927–1934.
- [20] C. Dimidov, G. Oriolo, and V. Trianni, "Random walks in swarm robotics: an experiment with kilobots," in *International Conference on Swarm Intelligence*. Springer, 2016, pp. 185–196.
- [21] J. Wilson, J. Timmis, and A. Tyrrell, "A hormone arbitration system for energy efficient foraging in robot swarms," in *Annual Conference Towards Autonomous Robotic Systems*. Springer, 2018, pp. 305–316.
- [22] J. Stradner, H. Hamann, T. Schmickl, and K. Crailsheim, "Analysis and implementation of an artificial homeostatic hormone system: A first case study in robotic hardware," in *IEEE/RSJ ICIRS*, 2009, pp. 595–600.
- [23] V. Trianni, E. Tuci, C. Ampatzis, and M. Dorigo, "Evolutionary swarm robotics: A theoretical and methodological itinerary from individual neurocontrollers to collective behaviors," *The Horizons of Evolutionary Robotics*, 2014, pp. 153–178.