

Research



Cite this article: Nauta J, Khaluf Y, Simoens P. 2020 Hybrid foraging in patchy environments using spatial memory. *J. R. Soc. Interface* **17**: 20200026.
<http://dx.doi.org/10.1098/rsif.2020.0026>

Received: 10 January 2020

Accepted: 21 April 2020

Subject Category:

Life Sciences—Physics interface

Subject Areas:

computational biology

Keywords:

Levy walk, random search, optimal foraging, spatial memory, heterogeneous environments

Author for correspondence:

Johannes Nauta

e-mail: johannes.nauta@ugent.be

Hybrid foraging in patchy environments using spatial memory

Johannes Nauta, Yara Khaluf and Pieter Simoens

Department of Information Technology—IDLab, Ghent University—imec, Technologiepark Zwijnaarde 126, 9052 Ghent, Belgium

JN, 0000-0002-5859-2729; YK, 0000-0002-5590-9321; PS, 0000-0002-9569-9373

Efficient random searches are essential to the survival of foragers searching for sparsely distributed targets. Lévy walks have been found to optimize the search over a wide range of constraints. When targets are distributed within patches, generating a spatial memory over the detected targets can be beneficial towards optimizing the search efficiency. Because foragers have limited memory, storing each target location separately is unrealistic. Instead, we propose incrementally learning a spatial distribution in favour of memorizing target locations. We demonstrate that an ensemble of Gaussian mixture models is a suitable candidate for such a spatial distribution. Using this, a hybrid foraging strategy is proposed, which interchanges random searches with informed movement. Informed movement results in displacements towards target locations, and is more likely to occur if the learned spatial distribution is correct. We show that, depending on the strength of the memory effects, foragers optimize search efficiencies by continuous revisitation of non-destructive targets. However, this negatively affects both the target and patch diversity, indicating that memory does not necessarily optimize multi-objective searches. Hence, the benefits of memory depend on the specific goals of the forager. Furthermore, through analysis of the distribution over walking distances of the forager, we show that memory changes the underlying walk characteristics. Specifically, the forager resorts to Brownian motion instead of Lévy walks, due to truncation of the long straight line displacements resulting from memory effects. This study provides a framework that opens up new avenues for investigating memory effects on foraging in sparse environments.

1. Introduction

Searching for targets whose distribution is unknown is an important problem in behavioural ecology. Locating scarce food sources is fundamental for the survival of the forager, and hence optimal foraging strategies must be employed in order to maximize survival probability [1,2]. Random searches can optimize the search for sparsely distributed targets [3–7], however including memory effects that change the random search into an informed search can highly influence the search efficiency of the forager [8–10]. Indeed, while random searches are a necessary component of any search strategy, its importance decays over time if the forager has the capacity to remember where targets are located. If all target locations are known, the random nature of the search can be completely ignored in favour of purely informed movement [11,12]. Thus, a trade-off between exploration (based on a non-informed random search) and exploitation (based on the gathered information) naturally arises when the target encounter rate is to be optimized through means other than random searches. This work proposes a hybrid foraging strategy based on memory, wherein random search and informed motion are alternated.

Although informed movement seems to enable further optimization of the search efficiency, any potential memory component still relies heavily on the efficiency of random searches for target localization. In particular, Lévy walks have been shown to optimize random searches in environments wherein

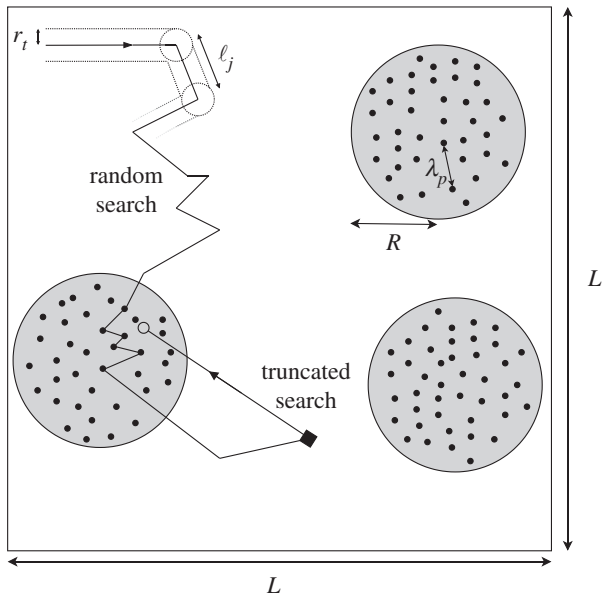


Figure 1. Illustrative example of the environment and the memory model for $N_p = 3$ patches with radius R and intrapatch mean free path λ_p . Targets (black dots) are distributed homogeneously (uniformly) within the patches, where each patch contains an equal amount of targets n_p . The random search is truncated (black square) according to the truncation probability (equation (2.2)) and the forager moves towards a sampled goal (open circle). After the informed motion, the random search continues, with the forager now back in the dense patch.

the target distribution is sparse [1,3,4,7,9,13], even in cases where targets were moving [14]. While search optimality of Lévy walks has been observed when targets are sparsely distributed, real-world environments are often both sparse and patchy [15–18], illustrating scale-free, fractal-like patterns [19–21]. These environments consist of patches containing densely distributed targets and interpatch regions where no targets are located (figure 1). Nonetheless, Lévy walks have been found to be robust to such fragmentation effects [9,22]. Moreover, Lévy walks have been found to encompass strategies which optimize patch diversity, wherein multiple-objectives play a key role in the random search process [13].

Patchy distributions within real-world environments remarkably alleviate the necessity for the foragers' cognitive abilities to remember the exact positions of targets. Specifically, in this study, we show that a forager can estimate patch locations (and thus the target locations) by learning a spatial distribution over detected targets. Furthermore, by sampling from the learned distribution, the forager can switch to informed movement in cases where no targets have been detected for some time. Hence, a hybrid foraging strategy emerges, in which random searches are alternated with bouts of informed motion as a result of memory induced effects. This improves target visitation rates by visiting regions that are known (or assumed by the model) to contain a high number of targets.

Spatial memory is less beneficial when foraging becomes destructive, meaning that targets are destroyed (e.g. consumed) upon detection. Obviously, learning a spatial distribution over target locations that constantly change due to destructive foraging is futile, since a detected target's location does not convey any information about possible future locations. Additionally, when the targets are constantly moving, for example in terrestrial prey–predator systems, the advantages of memory disappear. Indeed, as expected, memory has been found to be

more advantageous in predictable environments [23]. It is important to note that natural resources have the tendency to regenerate periodically, and hence a forager can, in principle, exploit knowledge of the regeneration times in order to return to patches rich in resources, e.g. seasonal changes. Such behaviour has been observed in Capuchin monkeys, which couple spatial locations to known (learned) time-varying patch qualities [24]. Stochastic regeneration rates are another type of regeneration observed in nature, for example the nectar in flowers [25], making learning the time dependence on the patch quality more difficult. While these types of interplay between resource regeneration and destructive foraging may correspond to more realistic searches, e.g. when describing foraging behaviour of marine predators that consume their targets [26], the empirically determined optimal Lévy parameter was often closer to the non-destructive optimum. Moreover, in ecology, patches are often dense [27], meaning that destructive foraging can effectively be approximated by non-destructive foraging behaviour since each target is within close proximity of another. This suggests that the consumption rate of such predators might be low enough, or that the regeneration rate within the patch is high enough, such that their foraging can be approximated as non-destructive. Hence, in the remainder of this study, we limit ourselves to non-destructive foraging tasks in static environments. Destructive foraging, foraging with known periodic or stochastic regeneration rates, and foraging for dynamic targets are topics for future research.

Next, we ask what features a good memory component for a forager in a patchy environment should have. An example of a simple random search that employs an infinite memory kernel is the extensively studied self-avoiding random walk. However, such models fail to capture memory effects that are useful for maximizing target visitation, since in non-destructive searches revisitation of targets is a very efficient foraging strategy. Other memory models that have been studied consist of the general direction wherein the forager prefers to move [28], location of a previously detected target [29,30] or a number of targets detected within the last set of steps [9]. However, more intricate memory models have been discussed based on diffusion models with drifts pointing towards patches with high target densities [10,31]. These works all indicate that memory is beneficial for maximizing target visitations, given that the foraging is non-destructive, regardless of the specific choice of memory model. Most interestingly, (time-varying) spatial memory has been experimentally verified in Capuchin monkeys [24,32,33]. Such results suggest that intelligent foragers are capable of maintaining a more complex memory model that is able to estimate distributions over targets.

In this work, we employ a Gaussian mixture model (GMM) that acts as the memory component, since it fits all requirements of a spatial memory. Due to their universal approximation properties [34,35], they can be used to learn a spatial target distribution over patchy landscapes from which the forager can sample goal locations. An additional advantage of a mixture of Gaussians is that each component can be weighted differently. This enables differentiating different patches, for example based on attributes of the patch, e.g. nutritional value. Furthermore, GMMs can be learned incrementally [36,37], allowing a forager with limited memory capacity to store positional target information within a small number of parameters that define the mixture model. Aside from real-world foragers not having infinite memory capabilities, a

major advantage is that such incremental learning procedures can adapt easily to possible changes in the target distribution. Actually, instead of learning a single GMM, we propose incremental learning of an ensemble of GMMs as to account for a model disagreement that measures the foragers' certainty of the learned distribution. This allows the foragers to make decisions based on the estimated correctness of their memory model [38–40]. More specifically, the random search component of the hybrid walk is more likely to be truncated when the forager is certain of its learned distribution over targets. We believe that the results presented in this paper are beneficial towards understanding the influence of memory in foraging in patchy environments.

The remainder of the paper is organized as follows. In §2, we describe and discuss the random search and the learning of the spatial distribution. We describe the patchy environment set-up and search efficiency measures in §3. In §4, we present the results of numerical simulations and discuss several properties of the resulting random search with informed movement. Finally, in §5, we conclude the paper and present final remarks and a future outlook on further experiments.

2. Model description

We start by specifying the informed search characteristics, which contains both a random search and a memory component. The forager alternates between the two components through truncation of the random search based on the learned memory model. Such a hybrid foraging strategy shows similarity with previously studied composite or adaptive Lévy walks [41–43], however, in this work, the change in strategy originates from the memory component of the informed search.

2.1. Informed search

Let us first discuss the random search component of the informed search. The random search is realized by assuming that foragers execute a Lévy random walk. Foragers following a Lévy walk pattern have their walk distances distributed according to a (truncated) power law $p(\ell) \sim \ell^{-\alpha}$, for $\ell_0 < \ell \leq L$ and where $1 < \alpha \leq 3$ is called the Lévy parameter. Note that $\alpha \leq 1$ corresponds to non-normalizable distributions. For $L \rightarrow \infty$, this power-law distribution of walk distances has the same asymptotic behaviour as the family of Lévy stable distributions [6,44], meaning that higher-order moments are infinite. However, in the real world such moments can never be infinite, since they correspond to infinitely large displacements. Therefore, we employ a truncated power-law distribution

$$p(\ell) = \begin{cases} \frac{\alpha-1}{\ell_0^{1-\alpha} - L^{1-\alpha}} \ell^{-\alpha} & \text{if } \ell_0 < \ell \leq L, \\ 0 & \text{otherwise,} \end{cases} \quad (2.1)$$

that effectively truncates walk distances at the environment size L . Here, ℓ_0 is the minimum travel distance, hence omitting steps of a much smaller scale that is irrelevant to the search process [45]. It is important to note convergence of the truncated power law of equation (2.1) towards a Gaussian process is ultraslow, meaning that the general characteristics of Lévy distributions are conserved during foraging [46].

In the asymptotic limit, the Lévy parameter defines the diffusion characteristics of the random walk, being

anomalous (superdiffusive) for $1 < \alpha \leq 3$ and normal (Brownian) for $\alpha \geq 3$ [47]. When $\alpha \rightarrow 1$, the forager moves ballistically, i.e. only straight-line motion is visible. For intermediate values, the random walk displays scale-free characteristics with statistically relevant large jumps that govern the diffusive capabilities of the forager. Such Lévy walks have been found to be optimal over a wide range of target distributions [22]. When foraging is non-destructive, $\alpha_{\text{opt}} \approx 2$ has been found to optimize the Lévy random search. In destructive foraging, ballistic motion with $\alpha_{\text{opt}} \rightarrow 1$ emerges as the optimal strategy. Intermediate searches, e.g. setting a degree of target revisitability through regeneration [48,49], give rise to intermediate values $1 \leq \alpha_{\text{opt}} \leq 2$ that optimize the random search. Additionally, values of $2 \leq \alpha_{\text{opt}} \leq 3$ can be found when a bias (e.g. an external drift, such as a current) is present [50]. For a more extensive overview of Lévy random searches, we refer the interested reader to previous works on Lévy walks in a foraging setting [6,27].

The actual informed search alternates the Lévy random search with memory induced informed motion, wherein the forager steers towards a goal location sampled from its spatial distribution. Foragers execute a Lévy random search that is truncated based on the learned model. Every j th walk step is followed according to the following rules:

- (a) Draw a walk distance ℓ_j from the (truncated) power-law distribution $p(\ell) \sim \ell^{-\alpha}$ (equation (2.1)), and an orientation angle θ sampled uniformly between 0 and 2π , and walk along the sampled path with fixed steps of size ℓ_0 .
- (b) While travelling the distance ℓ_j with steps of size ℓ_0 , the forager scans its direct environment within a radius r_t after each step. When a target is detected, the current walk is truncated. The forager moves to the target location¹ and stores the target location in its temporary (limited) memory. After detection of a target, a new walking distance and angle are sampled according to (a).
- (c) After each step, when no target is detected, the forager truncates the random search according to the number of steps L_0 wherein no target was encountered, a model disagreement φ (see §2.3) and a truncation parameter β (see below). The probability of truncation is computed at every step and is given by

$$p(L_0, \varphi) = 1 - \exp\left\{-\frac{\beta L_0}{\varphi}\right\}. \quad (2.2)$$

After the search has been truncated, the forager resorts to exploitation by sampling a goal state from its model $q_E(x, \Theta)$ (see equation (2.3))

$$x_{\text{goal}} \sim q_E(x|\Theta),$$

where Θ are the parameters of the learned spatial distribution. The forager then moves towards the sampled goal state. After this informed motion is finished, the forager starts the random search again at (a).

We fix the step size to the detection radius in equation (2.1), i.e. $\ell_0 = r_t$, hence naturally omitting steps smaller than the detection radius, since those are meaningless within a search context. Additionally, the upper truncation at L corresponds to the limit scale of the environment. Hence, the range $r_t < \ell \leq L$ appropriately defines the relevant scales of the search process.

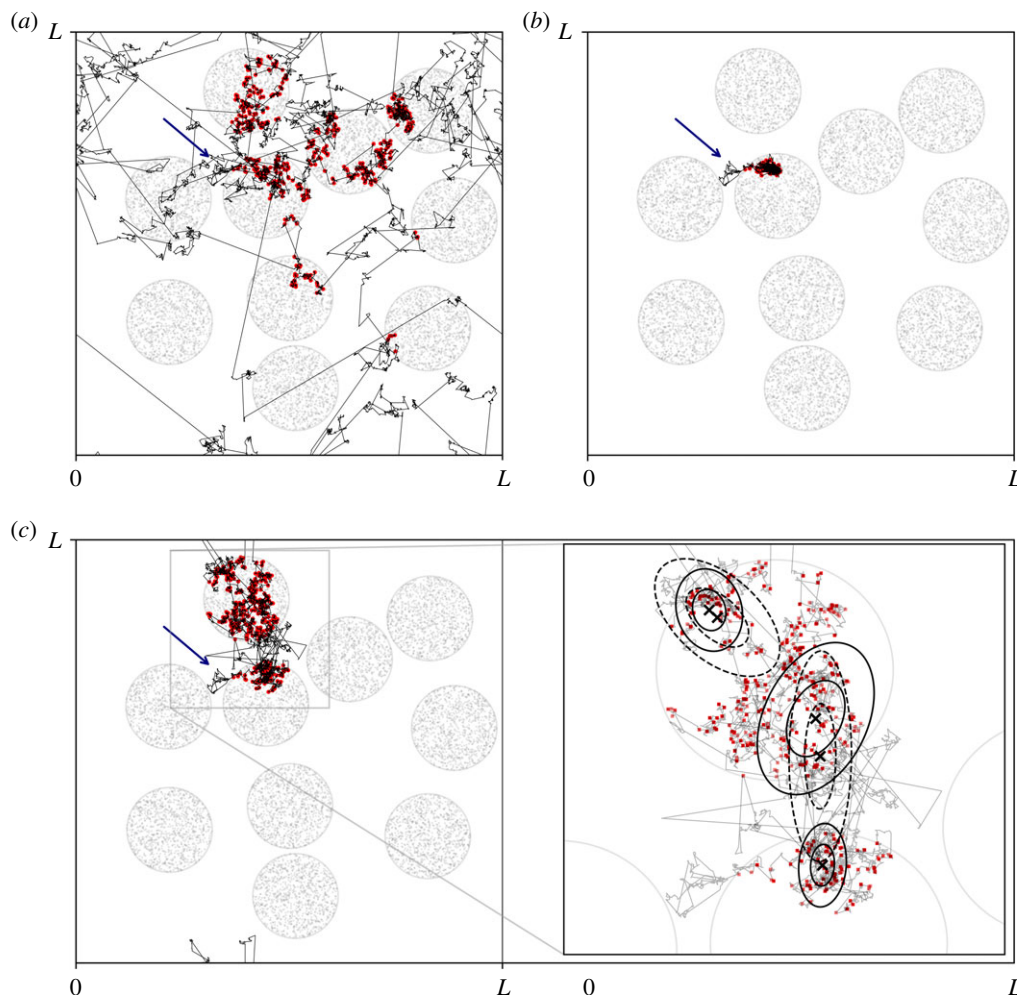


Figure 2. Examples of informed search for different truncation parameters β , where the task is to find 2×10^3 targets. Red dots indicate detected targets while the black lines are the trajectory of the forager. Patches and targets within the patches are denoted by the grey circles and dots respectively. The arrow indicates the (identical) starting position of the forager. (a) $\beta = 10^{-9}$: Naive forager that never employs a model, but continuously explores the environment using its random search. (b) $\beta = 10^2$: Greedy forager that almost always uses its model and hence quickly fixates around the first detected targets. (c) $\beta = 10^{-5}$: Intermediate forager that balances exploration and exploitation to generate a model over several patches. The inset displays the particular region in more detail, and includes continuous and dotted ellipses that represent variances of the components of the Gaussian mixtures of two ensemble members ($K = 2$), while crosses indicate the means.

First, let us discuss the truncation of the random search in more detail. Specifically, we adapt the framework of Zhao *et al.* [7] by including the model disagreement φ in the random search truncation probability of equation (2.2). This is indicated by the rule that *the longer the forager has travelled and the lower the model disagreement, the more likely the forager is to truncate its random search and exploit its knowledge*. The truncation parameter β defines the forager's behaviour (see figure 2 for an illustrative example). Low values of β correspond to *naive* foragers, which require an extremely low model disagreement before they trust their model enough to sample goal states from. This results in the forager relying on exploration, hence the random search will be the main component that determines the search efficiency. By contrast, relatively large values of β correspond to *greedy* foragers, that quickly resort to only sampling goal states from their model, even though their model might have a large disagreement and thus likely fails to describe the true target distribution. In between these two modes, there exists an *intermediate* forager that balances exploration and exploitation and hence is expected to optimize the search for targets. The truncation parameter β is thus an important metric that defines the behaviour of the forager.

The memory-induced informed movement after truncation of the random search is defined by taking steps along a calculated difference vector between the current position and the goal, again with steps of size $\ell_0 = r_t$. Thus, after the informed movement has finished, the agent is within a detection radius r_t of the sampled goal state x_{goal} . If the model is correct (i.e. it correctly models the sparse and patchy target distribution), each goal position corresponds to a position within a patch, and thus each goal state is, in principle, close to a (previously detected) target. However, when the model is incorrect this does not have to be true. During the informed movement, the forager continues to detect targets and truncates the informed motion once a target has been detected. This results in a more natural search, since in real scenarios it is unlikely that foragers will skip over available resources in favour of some arbitrary goal, given that we assume patches with the same attributes.

The target detection is based on a model of an effectively blind forager that only encounters targets within its perception range r_t [5]. While it can be argued whether animals capable of learning a spatial map have a limited perception radius, empirical results indicate that these animals still display super-diffusive random searches when foraging [29,51].

Furthermore, in this work, the sparsity is expressed in terms of the detection radius, hence increase of the detection radius does not change the outcome of the numerical experiments but simply changes the scale of the underlying foraging process. While many animals have evolved intricate perception methods, these are simply there to decrease the intrapatch target density λ_p , and hence such situations will be captured by discussing low values of the mean free path, which we do in §4.

Note that the truncation based on the model is sure to decrease the diffusivity of the resulting walk. In other words, the informed search might not carry the same diffusion characteristics of the free Lévy walk, which shows superdiffusive scaling behaviour. Depending on the target density, a switch to less diffusive searches might be realized by the forager, due to the searches revisiting a small set of targets within a patch with high target density, compliant with empirical data [29,33]. We both measure the characteristics of the resulting motion as well as introduce different search efficiency metrics in order to investigate the informed search in more detail.

2.2. Incremental learning procedure

In order to learn a spatial distribution over target locations, we employ incremental learning of a GMM with a variable number of mixture components. The variety in the number of components equips foragers with the ability to not overfit on the stored target locations, when a simpler model would suffice. Furthermore, we argue that the decision of the forager, e.g. the time at which to execute informed movement versus a random search, should be depending on the memory model. Hence, the GMM should be able to reflect some type of ‘certainty’ that determines the foragers’ belief that the model is accurate. If the model is indeed accurate, the forager should rely less on random search due to informed motion being the superior alternative. While a single GMM intrinsically captures the likelihood of the data given the model, ascribing ‘certainty’ to the model gives rise to an arbitrary threshold value that is difficult to determine. Instead, we deploy the foragers with an ensemble of GMMs from which the disagreement between each member of the ensemble and the full ensemble is able to reflect the model uncertainty [38–40]. Thus, the spatial distribution is given by an ensemble of K GMMs, each with M_k components (Gaussians):

$$q_E(x|\Theta) = \sum_{k=1}^K W_k q_k(x|\Theta^{(k)}), \quad (2.3)$$

$$\text{with } q_k(x|\Theta^{(k)}) = \sum_{j=1}^{M_k} w_j^{(k)} \mathcal{N}(\mu_j^{(k)}, \Sigma_j^{(k)}), \quad (2.4)$$

where $\Theta^{(k)} = \{(w_1^{(k)}, \mu_1^{(k)}, \Sigma_1^{(k)}), \dots, (w_{M_k}^{(k)}, \mu_{M_k}^{(k)}, \Sigma_{M_k}^{(k)})\}$ the parameters of the k th mixture model within the ensemble, $w_j^{(k)}$ the normalized weights of each multivariate Gaussian with mean $\mu_j^{(k)}$ and covariance $\Sigma_j^{(k)}$, and W_k the normalized ensemble weights. The normalization of the (ensemble) weights means that $\sum_j w_j^{(k)} = 1$ and $\sum_k W_k = 1$.

We assume that foragers have limited memory, hence each forager incrementally learns the GMM [37]. While the forager explores the environment and detects targets, it stores target locations in its memory up to a fixed maximum number N_{mem} . When the random search is truncated, the GMM is incrementally updated with the latest N_{mem} detected targets. Next,

the target locations are divided equally over the number of ensemble members K . It should be noted that a minimum number of data points is needed for the ensemble to be trained. Furthermore, GMMs are not suited to be computed over a single data point, hence we need at least a multiple of the total number of Gaussians as the minimum size of the dataset. After distribution of the data among the ensemble members, each member learns a new GMM over its subset of the recently detected N_{mem} target locations, with the number of mixture components between M_{min} and M_{max} . In other words, for each $M \in \{M_{\text{min}}, \dots, M_{\text{max}}\}$, a new GMM is learned. Then, compliant with existing statistical measures, the GMM with the lowest Bayesian information criterion (BIC) is selected as the new GMM [37]. Additionally, the weights of the new GMM are scaled by a forgetting factor f ($0 < f < 1$), which defines how much importance is attributed to the new, incoming data. Values $f \rightarrow 0$ indicate conservative foragers, which do not change their model with new incoming data, while $f \rightarrow 1$ represents progressive foragers that dispose of the current model in favour of the new one. Note that the forgetting factor additionally ensures normalization of the weights. Next, the previous GMM and the new GMM are joined to form the incrementally updated GMM. However, if the sum of components of the updated GMM exceeds the set maximum M_{max} , Gaussian components are merged. Which components are merged depends on their similarity, expressed in a symmetric KL divergence

$$D_{\text{KL}}^{\text{sym}} = \frac{1}{2}(D_{\text{KL}}(g_1, g_2) + D_{\text{KL}}(g_2, g_1)), \quad (2.5)$$

where $D_{\text{KL}}(g_i, g_j)$ the KL-divergence between two Gaussians $g_i \sim \mathcal{N}(\mu_i, \Sigma_i)$, given by

$$D_{\text{KL}}(g_i, g_j) = \log\left(\frac{\det \Sigma_j}{\det \Sigma_i}\right) + \text{Tr}(\Sigma_j^{-1} \Sigma_i) + (\mu_j - \mu_i)^T \Sigma_i^{-1} (\mu_j - \mu_i) - D, \quad (2.6)$$

where D is the number of dimensions of the Gaussian and $\det(A)$ and $\text{Tr}(A)$ the determinant respectively the trace of a matrix A . Components that carry the lowest values of the symmetric KL divergence are merged, i.e. components that are most similar, until the total number of components is equal to X_{max} . Merging of two Gaussians is defined by the following set of equations [36,37]:

$$w^* = w_1 + w_2 \quad (2.7)$$

$$\mu^* = \frac{w_1 \mu_1 + w_2 \mu_2}{w_1 + w_2} \quad (2.8)$$

$$\Sigma^* = \frac{w_1 \Sigma_1 + w_2 \Sigma_2}{w_1 + w_2} + \frac{w_1 w_2}{w_1 + w_2} (\mu_1 - \mu_2)(\mu_1 - \mu_2)^T, \quad (2.9)$$

where the asterisk denotes the new Gaussian. The merged Gaussians are deleted from the GMM and replaced by the newly merged Gaussian. An illustrative example of the final result of the learning procedure is depicted in figure 2c.

2.3. Truncation of the random search

With the spatial distribution in place, the forager needs to decide whether to truncate the current random search in favour of informed motion towards known (estimated) target locations. To compute the truncation probability of equation (2.2), the agent explicitly computes an uncertainty φ over the learned

Table 1. Overview of used variables.

environment	random search	model
L size of the environment	ℓ_0 minimum travel distance	φ model disagreement
R size of the patches	ℓ_j walk distance of walk step j	K number of ensemble members
N_p number of patches	α Lévy parameter	M_k number of Gaussian components
n_p number of targets per patch	η_s search efficiency	M_{\min} minimum number of components
N total number of targets	η_p patch search efficiency	M_{\max} maximum number of components
λ_e mean free path	p search truncation probability	
λ_p intrapatch mean free path	β truncation parameter	
r_t detection radius	L_0 steps without target detection	

ensemble. This uncertainty, the *model disagreement*, is defined as the KL divergence between each member of the ensemble and the full ensemble distribution

$$\varphi = \frac{1}{K} \sum_{k=1}^K D_{\text{KL}}[q_k(x|\Theta^{(k)})||q_E(x|\Theta)]. \quad (2.10)$$

Division by the number of ensemble members K ensures that the influence of memory is independent of the number of ensemble members. Since there does not exist a closed form solution for computing the KL-divergence between (ensemble) mixtures of Gaussians, like there is for two Gaussians in equation (2.6), we resort to Monte Carlo approximation [52], where

$$D_{\text{KL}}[q_k||q_E] = D_{\text{MC}}[q_k||q_E] = \frac{1}{N} \sum_{i=1}^N \log \frac{q_k(x_i|\Theta^{(k)})}{q_E(x_i|\Theta)}, \quad (2.11)$$

with $x_i \sim q_k(x_i|\Theta^{(k)})$ samples from ensemble member k . Note that even though the result is only exact in cases where $N \rightarrow \infty$, the forager does not require high accuracy of the integral, since it is not interested in the precision rather the order of magnitude of the disagreement.

2.4. Model overview

In summary, hybrid foraging encapsulates both a random search and informed motion based on the memory model. The forager learns an ensemble of GMMs over a dataset of target positions detected during the random search. From this ensemble, the model disagreement φ is computed, defining the uncertainty the forager has over the distribution of targets. The tuning parameter β determines the truncation probability $p(L_0, \varphi)$, which increases with the number of steps without detecting a target L_0 . The truncation probability depends on the model disagreement, where high model disagreements result in lower truncation probabilities, hence longer random searches, than when the model disagreement is low, corresponding to more frequent informed motion towards sampled goal targets. For an overview of all parameters used in this study see table 1.

3. Methods

3.1. Environment description

We consider a two-dimensional (2D) $L \times L$ space with periodic boundaries. The periodic boundary conditions effectively reflect infinite environments, akin to natural habitats being much larger than the forager itself. The forager is particle-like and is able to detect targets within a direct detection radius r_t . The task of the forager is to

find a fixed number of targets, i.e. the search is only halted after a specific number of targets have been detected. As an indication of the target sparsity, the mean free path λ_e can be computed by denoting that the cross-section of the targets equals $2r_t$ and the density $\rho = NL^{-2}$ can be expressed in the number of targets N and the area of the environment. The mean free path indicates the average distance between consecutive targets and is given by

$$\lambda_e = \frac{L^2}{2r_t N}. \quad (3.1)$$

Several sparsities can be studied by fixing the values for r_t and N , and varying λ_e which defines the environment size using equation (3.1) [4]. Note that the above expression for λ_e holds in general for any distribution over targets, thus including sparse and patchy distributions, however, it is less useful when the environment is patchy. In those cases, the *intrapatch* mean free path is more appropriate. In patchy environments, the targets are divided over N_p patches, with each patch containing n_p targets, hence $N = N_p n_p$. Then, the intrapatch mean free path λ_p (the mean free path within the patches) can be computed if we consider the patches to be non-overlapping circles with radius $R = wL < L$, $0 \leq w \leq 1$. Then we have

$$\lambda_p = \frac{\pi R^2}{2r_t n_p} = \frac{\pi w^2 L^2}{2r_t n_p}. \quad (3.2)$$

The intrapatch mean free path can be expressed in terms of the mean free path, through combining equations (3.1) and (3.2), giving

$$\lambda_p = \pi w^2 N_p \lambda_e. \quad (3.3)$$

Note that the mean free path and the intrapatch mean free path are indeed equal if the total area of all the distinct patches equals the total area of the environment. This allows us to forego the use of the mean free path λ_e in favour of the intrapatch mean free path λ_p , which is the relevant statistic in patchy environments.

In all following experiments, the forager is equipped with the task of finding a minimum of $N = 10^4$ targets. Within the environment, we distribute $n_p = 1000$ targets uniformly within each patch, where the radius of each patch is set to $R = 0.1L$. The total number of patches is set to $N_p = 10$. The size of the environment L is determined by the intrapatch mean free path λ_p , and is computed following equation (3.2). The step size, and subsequently the detection radius, is the unit for the system and is hence chosen to be $r_t = 1$. Results are averaged over 500 realizations, unless mentioned otherwise.

3.2. Search efficiency

In random searches, the efficiency of the process is often tightly interconnected with the survival of the forager, e.g. individuals might starve if they do not find enough food [53,54]. Hence, an appropriate metric that captures the efficiency is necessary.

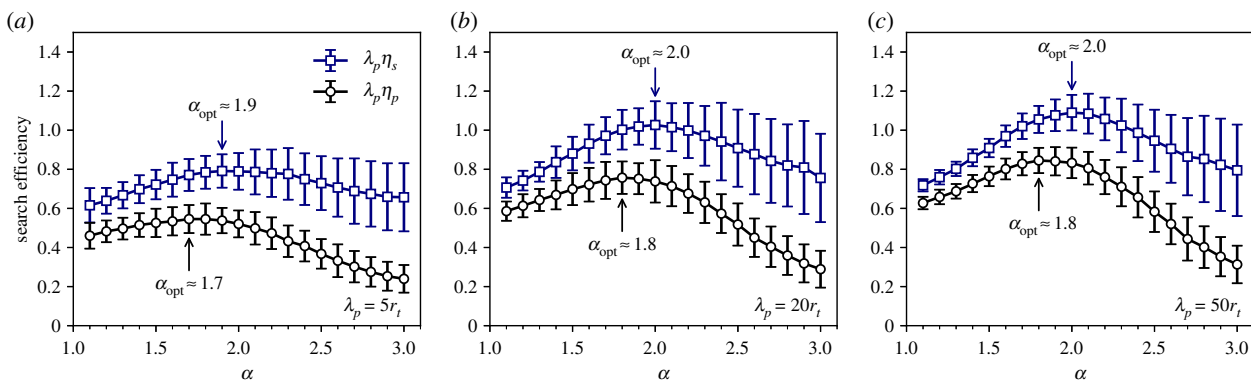


Figure 3. (a–c) Normalized (patch) search efficiency for $\beta = 0$, i.e. random search, for different Lévy walk parameters α . Blue squares represent the search efficiency $\eta_s \lambda_p$ (equation (3.4)), while black circles indicate the patch search efficiency $\eta_p \lambda_p$ (equation (3.5)). Optimal Lévy parameters are indicated with arrows. The search efficiency is always optimized for values $\alpha_{\text{opt}} \approx 2$, while the patch search efficiency is optimized for slightly lower values of α . Error bars represent 1 s.d.

Obviously, one cannot uniquely define a search efficiency metric that captures all possible constraints. Such constraints can vary from minimizing the time in between subsequent target visits [50] to minimizing energy consumption along the foraging trajectory [53–55].

We define two separate search efficiencies. The first search efficiency η_s defines the rate at which targets are detected. It is defined by

$$\eta_s = \frac{n}{d}, \quad (3.4)$$

where n indicates the total number of targets found and d denotes the total distance travelled by the forager. This definition of the search efficiency has been widely deployed for studying random searches [3–7,22]. However, as correctly identified in a previous study [13], it fails to capture patch diversity. While the search efficiency of equation (3.4) correctly measures the number of targets per travelled distance unit, it does not capture from which patches these targets originated. The result is that revisitation of a single target, or a few close-by targets within a single patch, greatly increases the search efficiency. We shall see that such behaviour indeed occurs when memory effects are strong, and hence the search efficiency that follows the above definition might not reflect more realistic scenarios, where patch diversity is also of importance. Multi-objective exploration might be beneficial for several reasons, all depending on the attributes of the patches (or targets). For example, the patches can differ in quality, hence constraining the search to a small subset of the number of patches might not optimize the total quality gain. Patches might also contain different types of targets, e.g. water and food, which are both necessarily for survival of the forager. Hence, following Wosniack *et al.* [13], we define the *patch search efficiency* as

$$\eta_p = \frac{1}{d} \frac{N_p^*}{N_p} \sum_{m=1}^{N_p^*} n_m \left(1 + \frac{|\bar{n} - n_m|}{n_m} \right)^{-1}, \quad (3.5)$$

where d is again the total distance travelled by the forager, n_m is the number of targets in patch m , $\bar{n} = M^{-1} \sum_m n_m$ the mean number of found targets per patch and N_p^* the number of distinct patches visited. Note that the denominator is minimized if $\bar{n} = n_m$, meaning that to optimize η_p , all patches should be visited an equal number of times. Mathematically, the search efficiency of equation (3.4) is recovered if $N_p^* = N_p$ and $\bar{n} = n_m$. This represents a homogeneous environment, i.e. each ‘patch’ only contains a single target. Furthermore, note that one can adapt the definition of equation (3.5) to account for the interpatch differences, by labelling targets by their different attributes

[13]. Extending both the memory model and the search efficiency to account for different types of targets is a topic for future work.

4. Results and discussion

4.1. Random search without memory ($\beta = 0$)

Let us first discuss the random search that occurs when the truncation probability equals 0, which corresponds to $\beta = 0$. In this case, there are no bouts of informed movement, hence the entire search is a random search. We study the random search efficiency in heterogeneous environments as a comparison against the proposed memory models, which shall be discussed in a later section. We measure the (patch) search efficiency of the foraging process for different values of the intrapatch mean free path λ_p . Results are shown in figure 3, and identify the existence of an optimal Lévy parameter α_{opt} . The search efficiency η_s is maximized for $\alpha_{\text{opt}} \approx 2$, regardless of the value of λ_p .

Recall that η_s is optimized through maximizing the number of targets detected while simultaneously minimizing the travelled distance. Intuitively, this is achieved by interchanging local search within dense patches (Brownian motion, $\alpha \rightarrow 3$) with global displacements in search of other regions with high target density (ballistic motion, $\alpha \rightarrow 1$). Sparse patches henceforth lead to intermediate values of $\alpha \approx 2$ being optimal. This is congruent with the well-established optimum in sparse, homogeneous environments [3,4,22].

By contrast, the patch search efficiency η_p is optimized by slightly more ballistic strategies, as indicated by the arrows that highlight the optimal value for α (figure 3). Recall that η_p is optimized when each patch is visited an equal number of times. Hence, due to the truncation of the walk at target detection, it is more difficult to exit a patch when patches are dense (small λ_p), while the forager needs to exit its current patch in order to maximize the patch search efficiency. The shift of the optimum to more ballistic strategies is therefore resulting from the fact that these foragers are more likely to exit the patch and thus are more likely to increase the number of distinct patches visited, hence the increase in search efficiency. Thus, when patch diversity is of key importance, more ballistic strategies are preferred (see also [13] and references therein).

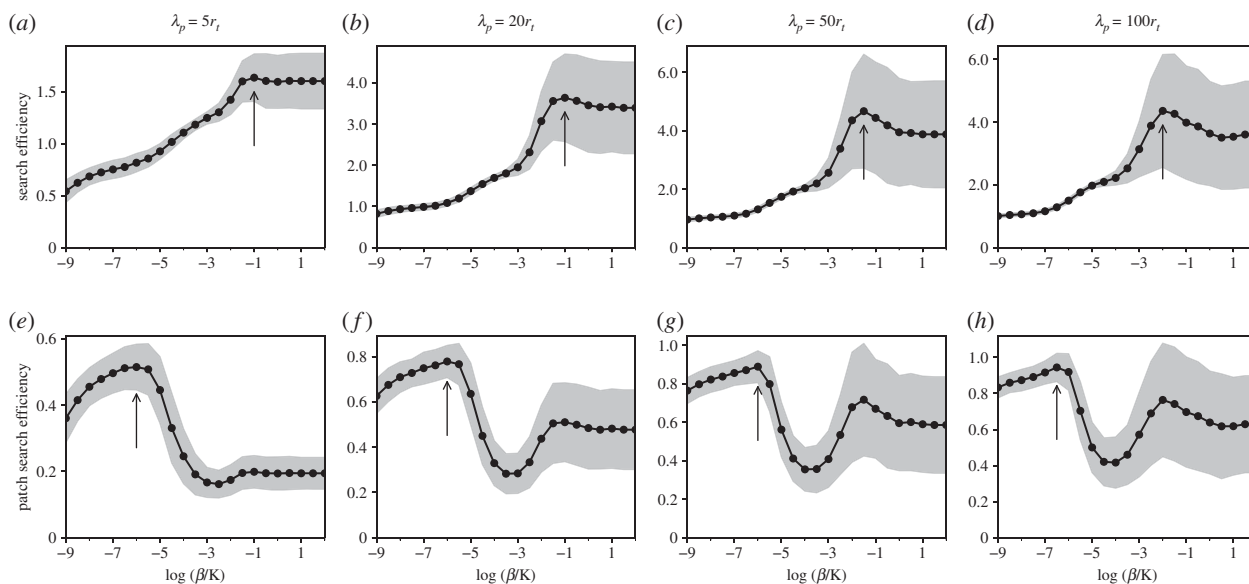


Figure 4. (a–h) Normalized search efficiency $\eta_s \lambda_p$ and patch search efficiency $\eta_p \lambda_p$ versus the scaled memory strength β/K for several intrapatch target densities λ_p . Low values of β/K correspond to highly explorative foragers, while larger values correspond to foragers that adopt their model quickly. The Lévy parameters are for the top row $\alpha = 2.0$, and for the bottom row $\alpha = 1.8$. Arrows indicate maxima, shaded areas represent 1 s.d.

4.2. Informed search with memory ($\beta > 0$)

Next, we study the influence of the truncation parameter β on the both search efficiencies. Note that the informed search is executed only after the model has been initialized, which occurs after the minimum number of targets $N_{\min} = 100$ have been detected through means of a random search. In other words, the informed search is always preceded by a random search, until a minimal number of targets have been detected through means of a random search. The targets found in this initial random search, act as the prior data for initializing the ensemble of GMMs that are incrementally updated with new batches of target positions found during the informed search. In all following experiments, we set the number of ensemble members $K = 3$, and the minimum and maximum number of components per member to be $1 \leq M_k \leq 10$. Note that the maximum number of components is equal to the number of patches N_p . For the Lévy parameter α , we chose the optimal value conform the random search, $\alpha = 2$ for the search efficiency η_s and $\alpha = 1.8$ for the patch search efficiency η_p . Results for various intrapatch mean free paths are shown in figure 4, which obviously indicate optima with respect to the memory strength β .

First, the search efficiency is maximized when memory effects are strong, figure 4a–d. As indicated earlier, revisitation of a select number of targets is highly beneficial for optimizing the search efficiency as defined in equation (3.4). Indeed, as illustrated in figure 5, the total number of unique targets (and patches) visited decreases when the memory strength increases. Hence, the forager indeed favours revisitation of a small fraction of the total number of targets, since its explorative motion is truncated due to strong memory effects, resulting in more informed movement towards the detected clusters of targets. Furthermore note that the search efficiency decreases when memory effects increase even more in strength. Since searches are ended after a fixed number of targets have been detected, the increase in variance results from an increase in variance of the distance travelled. This is indeed the case, as is shown in the inset of figure 6. Note that the variance increase is more substantial in very sparse environments, since target




Figure 5. Unique target and patch encounters versus the scaled memory strength β/K . (a) The percentage of unique targets visited, for several intrapatch mean free paths λ_p . The decrease indicates that as the memory strength increases, the forager favours revisitation of a select number of targets over more explorative behaviour that displays a wider variety of unique targets detected. Error bars represent 1 s.d. (b) The number of unique patches visited displays similar characteristics, furthermore indicating that foragers prefer to not venture out into interpatch space, so as to maximize the target encounter rate within a few patches found during the initial exploration (see also figure 2).

detection is inherently more difficult in those cases. Due to this sparsity, goals sampled from the spatial distribution learned over the targets detected during the initial random search may not be in close proximity of an actual target. Hence, the informed movement potentially has to be repeated several times, which results in the observed variance increase




Figure 6. The distance travelled versus the scaled memory strength β/K . The inset shows the standard deviation. More sparse distributions require larger distances to be travelled by the forager, and additionally increase the variance due to the fact that rates of detecting targets in very sparse environments are determined by chance (see text).

in the travel distance and subsequently both search efficiencies. Furthermore note that, as expected, more distance needs to be traversed to complete the search task whenever the target sparsity increases.

Next, the patch search efficiency displays interesting trends. It is optimized for much more explorative strategies, with β as low as 10^{-7} (conform to the explorative forager in figure 2). Thus, increased memory strength actually decreases the diversity of patch visitations, as also discussed above (figure 5), and more explorative behaviour is preferred. This originates from the choice of incremental learning of the ensemble of GMMs, since the ensemble members that were learned from the initial batch of detected targets can be sampled from much further into the informed search process. The influence of the initial set of detected targets can be adapted by the aforementioned forgetting rate, which controls how conservative the forager is when new targets are encountered. However, even when the forgetting rate is tuned such that all old target locations are forgotten, i.e. the incremental updating of the GMM is effectively replaced by learning a new GMM over the new target locations, the new target locations will most likely be around the old target locations when memory effects are strong. Hence, most of the different patches visited during the task are already visited during the initial random search, and memory does not increase the chance of detecting faraway regions rich in resources. In other words, foragers that strongly depend on their memory, are much more conservative and seldomly venture outwards to discover new patches, which is detrimental to patch search efficiency. If a high diversity of patches is preferred (or required, see §3.2), high memory strengths are sub-optimal. This indicates the long-term memory effects are not beneficial towards optimizing target diversity.

The patch search efficiency displays a minimum for intermediate memory strengths, which is more pronounced when targets are sparsely distributed within the patches (figure 4*e-h*). These minima arise from long bouts of exploration wherein no target was detected, only truncating the walk in favour of informed motion after a relatively large number of steps. These bouts result in the forager venturing out into unknown territories. However, due to the sparsity of the target distribution, it is unlikely to detect targets during these ventures. Hence, the forager travels more (unnecessary) distance before

giving up on the search and travelling back to known territories. Such hesitations, i.e. no true commitment to exploration, result in an increase in the travel distance that is not accompanied by an increase in the number of targets detected (as indicated in figure 5), hence the resulting decrease in patch search efficiency for intermediate values of β .

Finally, patch search efficiency again increases when memory strength increases, compared to intermediate memory strengths. This increase originates from the efficiency of the revisitation strategy of the forager, as the number of unique targets (and patches) visited do not change when the memory strength increases further (figure 5). However, venturing into unknown regions is suppressed due to higher truncation probabilities resulting from larger values of β . Essentially, this results in similar revisitation as intermediate memory strengths exhibited, while decreasing the distance travelled in between target revisitations (figure 6). In turn, this reduction in travel distance increases the patch search efficiency, however, the increase does not rise above the optimum reached for more explorative strategies, realized through small values of β .

4.3. Memory affects walk characteristics

Next, we discuss the characteristics of the walk that results from both the truncation at target detection as well as the memory. Specifically, the distribution over walk distances changes from a power law to an exponential distribution, which indeed does not possess the fat tails endemic to power-law distributions. As illustrated in the example trajectories of figure 2, memory greatly influences the random walk. Stronger memory effects (i.e. greedy foragers, high β), tend to more localized behaviour around the first few patches wherein targets are detected, hence disposing of the statistically relevant long-range displacements typical of the Lévy walk. Moreover, the power-law tail observed in Lévy walks is often absent when target densities are high, regardless of memory [1].

We demonstrate this in more detail by executing an informed search for a fixed number of targets, wherein we define an *episode* to be a time window wherein the model remains fixed. Thus, an episode ends at the same time the random search is truncated and the model is (incrementally) updated. First, the model disagreement φ is plotted, including the single-step truncation probability $p(1, \varphi)$, in figure 7*a,b*. The model disagreement decreases as the number of episodes (model updates) increases. Naturally, the corresponding truncation probability depends heavily on β , as the intermediate forager has a single step truncation probability of effectively 0 (figure 7*a*). However, the greedy forager is very likely to truncate its random search after a single step after a few model updates due to the combination of a low model disagreement φ and high value of β . As a result, the greedy forager repetitively only visits targets from memory and omits exploration entirely.

Next, we record the walk distances within each episode, which can be truncated due to target encounters. Using the distribution over walk distances within each episode, we can fit a (truncated) power-law distribution and compute an approximate value for the Lévy parameter α . However, fitting a distribution is misleading due to the fact that the underlying data might not be described by the candidate function [56–59]. Therefore, the log-likelihood ratio (LLR) between a (truncated) power law, which indicates anomalous diffusion, and an exponential distribution, which is indicative of Brownian-like

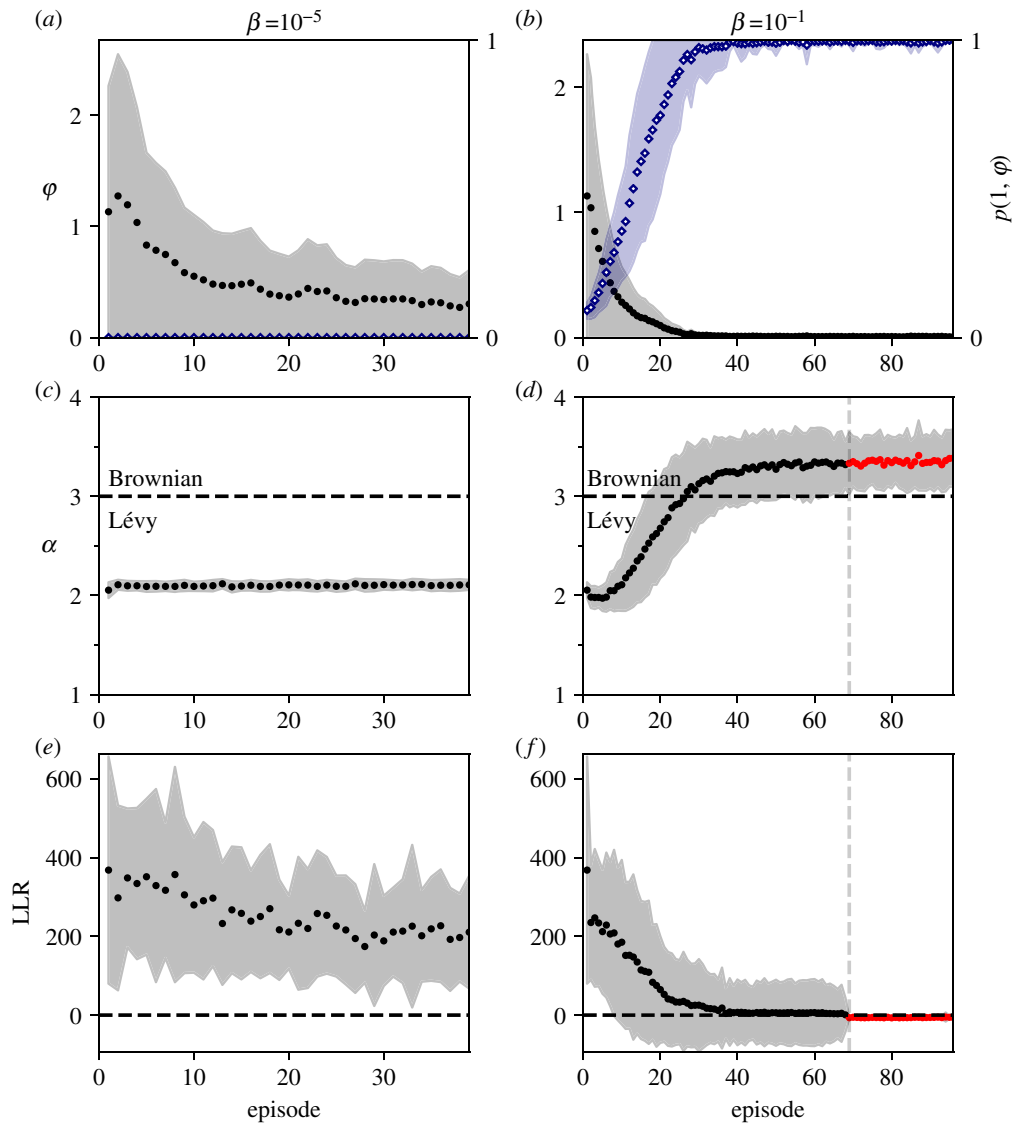


Figure 7. (a–f) Influence of memory on the power-law distribution over walk distances for foragers searching with Lévy parameter $\alpha = 2$ in an environment with $\lambda_p = 20r_t$. Two types of foragers—intermediate ($\beta = 10^{-5}$, left), and greedy ($\beta = 10^{-1}$, right)—have executed an informed search to search for 10^4 targets. (Top) The average model disagreement φ continuously decreases as the model is updated. The truncation probability $p(1, \varphi)$ is shown in blue, with $L_0 = 1$. The approximate power-law exponent (middle) is computed from the distribution over walk distances, with the corresponding log-likelihood ratio (bottom) between a power-law distribution and an exponential distribution. Each data point corresponds to a single episode (see text). Red points display episodes for which the average LLR < 0 , while the vertical dotted line highlights the first occurrence. (a) + (c) + (e) display an intermediate forager that displays a power-law distribution across the full task. (b) + (d) + (f) display a greedy forager, which changes its walk distance distribution to an exponential (Brownian) one as its model is incrementally updated. Shaded areas represent 1 s.d., averaged over 50 different realizations.

diffusion, is computed. For this computation, we use existing and well-established libraries [57].

We report results for two distinct types of foragers; intermediate and greedy (figure 2). The naive forager is not shown, since those do not adapt their model due to their truncation probability being effectively zero, hence the search consists of a single episode which indeed obeys a power-law distribution over walk distances (results not shown). As seen in figure 7, the intermediate forager has walk distance distributions that obey a power law, up until the task is completed. In contrast, the greedy forager changes its walk behaviour to Brownian motion the more its spatial distribution is updated. The result is that, approximately, both the fitted power-law component $\alpha \geq 3$, and $\text{LLR} < 0$, indicating that the distribution over walk distances does not follow a power law but is more likely to obey an exponential distribution, i.e. Brownian motion [1]. Memory thus affects

the walk distance distribution, inducing less diffusive motion in favour of revisitation of a select area rich in targets. This decrease in diffusivity directly results in the previously mentioned suboptimal patch search efficiencies reached when memory effects are strong, since those are optimized with more ballistic strategies (see figure 3 and [13]).

Whereas assigning power-law distributions to empirical datasets of walk distances needs to be handled with caution [58,59], our rudimentary study of the underlying distributions illustrates that the diffusivity is not necessarily determined by the underlying sampling procedure [1]. Indeed, random walks with memory induced relocations to previously visited locations, have been shown to display subdiffusion [30]. Moreover, Brownian motion induced by memory is favoured over more ballistic Lévy walks when targets are more difficult to encounter [10], indicating that the switch to more Brownian-like motion aids non-destructive

target detection in very sparse environments. Additional careful analysis on the effect of spatial memory on the movement pattern of the forager is beyond the scope of this paper.

5. Conclusion

We have investigated the effects of spatial memory on the search efficiency in patchy environments. Spatial memory constituted an ensemble of GMM. The resulting hybrid foraging strategy, alternates Lévy walks as random searches with memory induced walks. The truncation probability, which indicates the switch between the two distinct walks, increases with the number of steps since the last target detection and the decrease of the model disagreement. We have shown that non-destructive foraging is optimized when memory effects are strong and the search efficiency is defined solely through the rate of target detection per unit travelled. However, strong memory effects result in high revisitation rates of a select fraction of the targets, often within a small portion of the available patches. Thus, explorative behaviour is penalized due to increases in the travelled distance not being accompanied by similar increases in the number of targets detected. This identifies a trade-off between the search efficiency and patch diversity. The trade-off is not trivially solved by intermediate memory strengths, since in sparse environments the patch search efficiency is actually minimized when the memory strength is intermediate. As a result, the benefits of memory heavily depends on the needs of the forager. If a high diversity of targets is required for survival, then more explorative motion is preferred, whereas memory serves more useful in sparse environments in which targets can be revisited indefinitely.

Where we have assumed an infinite regeneration rate by assuming non-destructive foraging, we believe that the results might also be beneficial for finite regeneration rates. Notably, the benefits the forager obtains from using a spatial memory strongly depend on the regeneration rates of targets. When resources are sparsely distributed, but do regenerate over potentially long periods of time, memory might serve as a very useful tool that increases survival chance due to enabling revisitation of previously exhausted patches that have regenerated [60]. Moreover, finely tuning memory strengths with potential periodicity of the availability of the targets, e.g. seasonal growth of food, might greatly reduce search times and energy consumption and as a result increase the search efficiency [61]. In addition, random walks without memory also exhibit a switch from superdiffusive walks to ballistic motion as the target regeneration time changes from zero (non-destructive) to infinity (destructive) [49].

Furthermore using an ensemble of GMMs as a spatial distribution enables researching attribute dependent behaviour. Since each member of the ensemble consists of a mixture of Gaussians, weighting the Gaussians differently when

sampling goal states can be coupled to certain attributes of the targets wherever the Gaussian was fitted, e.g. patch or target quality. Additionally, each member can represent different types of targets in order to account for balanced target detections, furthermore accompanied by the aforementioned possible adaptations of the patch search efficiency. Future studies might indicate that foragers that aim to optimize overall target diversity, benefit more from an ensemble of differently weighted spatial distributions.

Lastly, it is important to note that the Lévy walk paradigm has been criticized as being unrealistic [45], mainly because traditional implementations of the Lévy walk omit directionality resulting from a continuous decision process. An effect of the continuity of the foraging process, is that any organism can change its behaviour at any instant in time [62]. In this work, the decision process of the forager is continuous as the forager decides to continue or truncate the search based on its current available information from both from its perception as well as its memory. An additional critique was that the spatial scales wherever the walk characteristics are determined should be relevant to the foraging process, something that we have achieved by sampling walk distances using a truncated power law. Hence, while further investigation on the motion of foragers is warranted, we have directly tackled the main criticisms of the Lévy walk paradigm by inclusion of the truncation probability depending on memory.

This study has indicated that spatial memory is not necessarily beneficial towards detecting a diverse set of targets. We have identified strengths and weaknesses of a potential spatial memory candidate, which suggests that the balance between using memory versus naive random walks is fickle and is in much need of further study.

Data accessibility. This article has no additional data.

Authors' contributions. J.N. conceived of the study, designed the study, carried out the experiments, carried out data analysis and drafted the manuscript. Y.K. conceived of the study, coordinated the study and critically revised the manuscript. P.S. coordinated the study and critically revised the manuscript. All authors gave final approval for publication and agree to be held accountable for the work performed therein.

Competing interests. We declare we have no competing interest.

Funding. No funding has been received for this article.

Acknowledgements. We thank Marina E. Wosniack for very useful discussions.

Endnotes

¹In our study, the perception radius is of equal size as the step size of the forager, trivializing the process of the forager moving to the target location. However, when the perception radius is much larger than the step size, the forager should compute a relative angle between itself and the detected target and travel in that direction until it is on the target.

References

1. Wosniack ME, Santos MC, Raposo EP, Viswanathan GM. 2017 The evolutionary origins of Lévy walk foraging. *PLoS Comput. Biol.* **13**, e1005774. (doi:10.1371/journal.pcbi.1005774)
2. Anselme P, Otto T, Güntürkün O. 2018 Foraging motivation favors the occurrence of Lévy walks. *Behav. Processes* **147**, 48–60. (doi:10.1016/j.beproc.2017.12.014)
3. Viswanathan GM, Buldyrev SV, Havlin S, Da Luz M, Raposo E, Stanley HE. 1999 Optimizing the success of random searches. *Nature* **401**, 911. (doi:10.1038/44831)

4. Bartumeus F, Viswanathan GM, Catalan J. 2005 Animal search strategies: a quantitative random-walk analysis. *Ecology* **86**, 3078–3087. (doi:10.1890/04-1806)
5. Reynolds A. 2015 Extending Lévy search theory from one to higher dimensions: Lévy walking favours the blind. *Proc. R. Soc. A* **471**, 20150123. (doi:10.1098/rspa.2015.0123)
6. Zaburdaev V, Denisov S, Klafter J. 2015 Lévy walks. *Rev. Mod. Phys.* **87**, 483. (doi:10.1103/RevModPhys.87.483)
7. Zhao K, Jurdak R, Liu J, Westcott D, Kusy B, Parry H, Sommer P, McKeown A. 2015 Optimal Lévy-flight foraging in a finite landscape. *J. R. Soc. Interface* **12**, 20141158. (doi:10.1098/rsif.2014.1158)
8. Olsson O, S Brown J. 2006 The foraging benefits of information and the penalty of ignorance. *Oikos* **112**, 260–273. (doi:10.1111/j.0030-1299.2006.13548.x)
9. Ferreira A, Raposo E, Viswanathan G, Da Luz M. 2012 The influence of the environment on Lévy random search efficiency: fractality and memory effects. *Physica A* **391**, 3234–3246. (doi:10.1016/j.physa.2012.01.028)
10. Bracis C, Gurarie E, Van Moorter B, Goodwin RA. 2015 Memory effects on movement behavior in animal foraging. *PLoS ONE* **10**, e0136057. (doi:10.1371/journal.pone.0136057)
11. Menzel R *et al.* 2005 Honey bees navigate according to a map-like spatial memory. *Proc. Natl Acad. Sci. USA* **102**, 3040–3045. (doi:10.1073/pnas.0408550102)
12. Fagan WF *et al.* 2013 Spatial memory and animal movement. *Ecol. Lett.* **16**, 1316–1329. (doi:10.1111/ele.12165)
13. Wosniack M, Raposo E, Viswanathan G. 2015 Efficient search of multiple types of targets. *Phys. Rev. E* **92**, 062135. (doi:10.1103/PhysRevE.92.062135)
14. Bartumeus F, Catalan J, Fulco UL, Lyra ML, Viswanathan GM. 2002 Optimizing the encounter rate in biological interactions: Lévy versus Brownian strategies. *Phys. Rev. Lett.* **88**, 097901. (doi:10.1103/PhysRevLett.88.097901)
15. Levin SA. 2000 Multiple scales and the maintenance of biodiversity. *Ecosystems* **3**, 498–506. (doi:10.1007/s100210000044)
16. Weimerskirch H. 2007 Are seabirds foraging for unpredictable resources? *Deep Sea Res. Part II* **54**, 211–223. (doi:10.1016/j.dsr2.2006.11.013)
17. Kéfi S, Rietkerk M, Alados CL, Pueyo Y, Papanastasis VP, ElAich A, De Ruiter PC. 2007 Spatial vegetation patterns and imminent desertification in Mediterranean arid ecosystems. *Nature* **449**, 213. (doi:10.1038/nature06111)
18. Khaluf Y, Ferrante E, Simoens P, Huepe C. 2017 Scale invariance in natural and artificial collective systems: a review. *J. R. Soc. Interface* **14**, 20170662. (doi:10.1098/rsif.2017.0662)
19. Fauchald P. 1999 Foraging in a hierarchical patch system. *Am. Nat.* **153**, 603–613. (doi:10.1086/303203)
20. Johnson CJ, Boyce MS, Mulders R, Gunn A, Gau RJ, Cluff HD, Case RL. 2004 Quantifying patch distribution at multiple spatial scales: applications to wildlife–habitat models. *Landscape Ecol.* **19**, 869–882. (doi:10.1007/s10980-004-0246-7)
21. Wheatley M. 2010 Domains of scale in forest-landscape metrics: implications for species–habitat modeling. *Acta Oecol.* **36**, 259–267. (doi:10.1016/j.actao.2009.12.003)
22. Wosniack ME, Santos MC, Raposo EP, Viswanathan GM. 2015 Robustness of optimal random searches in fragmented environments. *Phys. Rev. E* **91**, 052119. (doi:10.1103/PhysRevE.91.052119)
23. Mueller T, Fagan WF, Grimm V. 2011 Integrating individual search and navigation behaviors in mechanistic movement models. *Theor. Ecol.* **4**, 341–355. (doi:10.1007/s12080-010-0081-1)
24. Tujague MP, Janson CH. 2017 Wild capuchin monkeys anticipate the amount of ripe fruit in natural trees. *Anim. Cogn.* **20**, 841–853. (doi:10.1007/s10071-017-1105-7)
25. Possingham HP. 1988 A model of resource renewal and depletion: applications to the distribution and abundance of nectar in flowers. *Theor. Popul. Biol.* **33**, 138–160. (doi:10.1016/0040-5809(88)90010-X)
26. Sims DW *et al.* 2008 Scaling laws of marine predator search behaviour. *Nature* **451**, 1098. (doi:10.1038/nature06518)
27. Viswanathan GM, Da Luz MG, Raposo EP, Stanley HE. 2011 *The physics of foraging: an introduction to random searches and biological encounters*. Cambridge, UK: Cambridge University Press.
28. Schütz GM, Trimper S. 2004 Elephants can always remember: exact long-range memory effects in a non-Markovian random walk. *Phys. Rev. E* **70**, 045101. (doi:10.1103/PhysRevE.70.045101)
29. Boyer D, Solis-Salas C. 2014 Random walks with preferential relocations to places visited in the past and their application to biology. *Phys. Rev. Lett.* **112**, 240601. (doi:10.1103/PhysRevLett.112.240601)
30. Campos D, Méndez Vmc. 2019 Recurrence time correlations in random walks with preferential relocation to visited places. *Phys. Rev. E* **99**, 062137. (doi:10.1103/PhysRevE.99.062137)
31. Calhoun AJ, Chalasani SH, Sharpee TO. 2014 Maximally informative foraging by *Caenorhabditis elegans*. *Elife* **3**, e04220. (doi:10.7554/eLife.04220)
32. Janson CH. 1998 Experimental evidence for spatial memory in foraging wild capuchin monkeys, *Cebus apella*. *Anim. Behav.* **55**, 1229–1243. (doi:10.1006/anbe.1997.0688)
33. Boyer D, Crofoot MC, Walsh PD. 2011 Non-random walks in monkeys and humans. *J. R. Soc. Interface* **9**, 842–847. (doi:10.1098/rsif.2011.0582)
34. Goodfellow I, Bengio Y, Courville A. 2016 *Deep learning*. Cambridge, MA: MIT Press.
35. Plataniotis KN, Hatzinakos D. 2017 Gaussian mixtures and their applications to signal processing. In *Advanced signal processing handbook*, pp. 89–124. Boca Raton, FL: CRC Press.
36. Bouchachia A, Vanaret C. 2011 Incremental learning based on growing gaussian mixture models. In *2011 10th Int. Conf. on Machine Learning and Applications and Workshops*. vol. 2, pp. 47–52. Piscataway, NJ: IEEE.
37. Acevedo Valle JM, Trejo Ramírez KA, Angulo Bah'ón C. 2017 Multivariate regression with incremental learning of Gaussian mixture models. In *Recent Advances in Artificial Intelligence Research and Development: Proc. of the 20th Int. Conf. of the Catalan Association for Artificial Intelligence, Deltebre, Terres de l'Ebre, Spain, 25–27 October*, pp. 196–205. Amsterdam, The Netherlands: IOS Press.
38. Glodek M, Schels M, Schwenker F. 2013 Ensemble Gaussian mixture models for probability density estimation. *Comput. Stat.* **28**, 127–138. (doi:10.1007/s00180-012-0374-5)
39. Lakshminarayanan B, Pritzel A, Blundell C. 2017 Simple and scalable predictive uncertainty estimation using deep ensembles. In *Advances in Neural Information Processing Systems 30: 31st Ann. Conf. on Neural Information Processing Systems (NIPS 2017), 4–9 December 2017, Long Beach, CA*, pp. 6402–6143.
40. Chua K, Calandra R, McAllister R, Levine S. 2018 Deep reinforcement learning in a handful of trials using probabilistic dynamics models. In *Advances in Neural Information Processing Systems 31: 32nd Ann. Conf. on Neural Information Processing Systems 2018, NeurIPS 2018, 3–8 December, Montréal, Canada*, pp. 4754–4765.
41. Plank M, James A. 2008 Optimal foraging: Lévy pattern or process? *J. R. Soc. Interface* **5**, 1077–1086. (doi:10.1098/rsif.2008.0006)
42. Reynolds A. 2009 Adaptive Lévy walks can outperform composite Brownian walks in non-destructive random searching scenarios. *Physica A* **388**, 561–564. (doi:10.1016/j.physa.2008.11.007)
43. Nolting BC, Hinkelman TM, Brassil CE, Tenhumberg B. 2015 Composite random search strategies based on non-directional sensory cues. *Ecol. Complexity* **22**, 126–138. (doi:10.1016/j.ecocom.2015.03.002)
44. Klafter J, Zumofen G, Schlesinger M. 1995 Lévy description of anomalous diffusion in dynamical systems. In *Lévy flights and related topics in physics*, pp. 196–215. New York, NY: Springer.
45. Pyke GH. 2015 Understanding movements of organisms: it's time to abandon the Lévy foraging hypothesis. *Methods Ecol. Evol.* **6**, 1–16. (doi:10.1111/2041-210X.12298)
46. Mantegna RN, Stanley HE. 1994 Stochastic process with ultraslow convergence to a Gaussian: the truncated Lévy flight. *Phys. Rev. Lett.* **73**, 2946–2949. (doi:10.1103/PhysRevLett.73.2946)
47. Bouchaud JP, Georges A. 1990 Anomalous diffusion in disordered media: statistical mechanisms, models and physical applications. *Phys. Rep.* **195**, 127–293. (doi:10.1016/0370-1573(90)90099-N)
48. Raposo EP, Buldyrev SV, Santos MC, Stanley HE, Viswanathan GM. 2003 Dynamical robustness of Lévy search strategies. *Phys. Rev. Lett.* **91**, 240601. (doi:10.1103/PhysRevLett.91.240601)
49. Santos M, Raposo E, Viswanathan G, Da Luz M. 2004 Optimal random searches of revisitable targets: crossover from superdiffusive to ballistic

- random walks. *EPL (Europhysics Letters)* **67**, 734. (doi:10.1209/epl/i2004-10114-9)
50. Palyulin VV, Chechkin AV, Metzler R. 2014 Lévy flights do not always optimize random blind search for sparse targets. *Proc. Natl Acad. Sci. USA* **111**, 2931–2936. (doi:10.1073/pnas.1320424111)
 51. Ramos-Fernández G, Mateos JL, Miramontes O, Cocho G, Larralde H, Ayala-Orozco B. 2004 Lévy walk patterns in the foraging movements of spider monkeys (*Ateles geoffroyi*). *Behav. Ecol. Sociobiol.* **55**, 223–230. (doi:10.1007/s00265-003-0700-6)
 52. Hershey JR, Olsen PA. 2007 Approximating the Kullback Leibler divergence between Gaussian mixture models. In *2007 IEEE Int. Conf. on Acoustics, Speech and Signal Processing-ICASSP'07*, vol. 4, p. IV–317. Piscataway, NJ: IEEE.
 53. Faustino C, Da Silva L, Da Luz M, Raposo E, Viswanathan G. 2007 Search dynamics at the edge of extinction: anomalous diffusion as a critical survival state. *EPL (Europhysics Letters)* **77**, 30002. (doi:10.1209/0295-5075/77/30002)
 54. Faustino C, Lyra M, Raposo E, Viswanathan G. 2012 The universality class of random searches in critically scarce environments. *EPL (Europhysics Letters)* **97**, 50005. (doi:10.1209/0295-5075/97/50005)
 55. Ribeiro-Neto P, Raposo E, Araújo H, Faustino C, Viswanathan G. 2012 Dissipative Lévy random searches: universal behavior at low target density. *Phys. Rev. E* **86**, 061102. (doi:10.1103/PhysRevE.86.061102)
 56. Clauset A, Shalizi CR, Newman MEJ. 2009 Power-law distributions in empirical data. *SIAM Rev.* **51**, 661–703. (doi:10.1137/070710111)
 57. Alstott J, Bullmore E, Plenz D. 2014 Powerlaw: a Python package for analysis of heavy-tailed distributions. *PLoS ONE* **9**, e85777. (doi:10.1371/journal.pone.0085777)
 58. Broido AD, Clauset A. 2019 Scale-free networks are rare. *Nat. Commun.* **10**, 1017. (doi:10.1038/s41467-019-08746-5)
 59. Rausch I, Khaluf Y, Simoens P. 2019 Scale-free features in collective robot foraging. *Appl. Sci.* **9**, 2667. (doi:10.3390/app9132667)
 60. Winter Y, Stich KP. 2005 Foraging in a complex naturalistic environment: capacity of spatial working memory in flower bats. *J. Exp. Biol.* **208**, 539–548. (doi:10.1242/jeb.01416)
 61. Louise RL, Simon B, Christophe B, Simon CJ. 2017 Spatial memory shapes density dependence in population dynamics. *Proc. R. Soc. B* **284**, 20171411. (doi:10.1098/rspb.2017.1411)
 62. McFarland D. 2017 Decision making in animals. *Nature* **269**, 15–21. (doi:10.1038/269015a0)