Evolutionary foraging in grids: Intermittent search dynamics emerge in finite, depletable landscapes

Shailendra Bhandari
Department of Computer Science
OsloMet – Oslo Metropolitan University
St. Olavs plass, N-0130 Oslo, Norway
shailendra.bhandari@oslomet.no

Alex Szorkovszky
Simula Research Laboratory
Numerical Analysis and Scientific Computing
Oslo, 0164, Norway
alex@simula.no

Anis Yazidi
Department of Informatics
University of Oslo
Gaustadalléen 23B, 0373 Oslo, Norway
Department of Computer Science
OsloMet – Oslo Metropolitan University
St. Olavs plass, N-0130 Oslo, Norway
anisy@uio.no

Pedro G. Lind
Department of Technology
Kristiania University of Applied Sciences
Oslo, Norway
Department of Computer Science
OsloMet – Oslo Metropolitan University
St. Olavs plass, N-0130 Oslo, Norway
pedrolin@oslomet.no

Abstract

How search strategies evolve in finite, depletable landscapes remains a central question in foraging theory. We study this problem with an evolutionary simulation in which agents forage on a two-dimensional toroidal lattice containing non-renewable resources distributed either uniformly or as Lévy dust. Each agent carries a heritable genome encoding step lengths, velocities, and turning angles, and selection acts on a fitness function combining energetic gain, movement cost, and coverage efficiency. By allowing movement traits to evolve without imposing a prescribed power-law step-length distribution, we test whether the evolved trajectories are better described by intermittent-search or Lévy-walk dynamics. Strikingly, our results indicate that, in the finite depletion-driven landscapes considered here, evolved search is more consistent with intermittent dynamics than with strict scale-free Lévy motion. To characterize the effective dynamics of the evolved trajectories, we take the average adjusted coefficient of determination when fitting second- and fourth-order displacement moments to intermittent search and Lévy-walk models. While a Lévy-like random walk fits with the numerical results from the evolutionary search very well (mean adjusted $\bar{R}^2 > 0.9$ in most tested conditions), the intermittent search achieves a closer fit, with mean adjusted coefficients of determination ($\bar{R}^2 > 0.99$) for all tested resource distributions. This preference holds across the tested grid sizes and resource densities. Five independent evolutionary runs per environment on a 503×503 grid at nominal resource density $\rho = 0.15$ reproduce this preference for the uniform environment and the five Lévy-dust environments. Evolution rapidly reshapes the movement genome toward short displacements while retaining a sparse tail of longer relocations, consistent with local exploitation punctuated by occasional transfer. The framework provides a controlled setting for studying how search rules emerge under resource limitation and may inform resource-constrained exploration in autonomous systems. Project webpage: https://evo-foraging.github.io/.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

1 Introduction

Foraging efficiency is a central problem in foraging theory, because survival and reproductive success often depend on how organisms balance resource gain against movement cost in uncertain environments [1, 2, 3, 4]. A major line of work has focused on the Lévy walk (LW) foraging hypothesis, which proposes that power-law-distributed step lengths can optimize search in sparse landscapes [5, 6, 7]. The inverse-square case $p(s) \propto s^{-2}$, often referred to as the Cauchy walk, is a particularly important benchmark because it corresponds to the classical exponent near 2 associated with efficient Lévy foraging [8]. Empirical and theoretical studies have reported Lévy-like movement across many biological systems, including marine predators, birds, insects, and human search behavior [9, 10, 11, 12]. Lévy-like movement has also influenced algorithmic search and optimization, where rare long relocations can improve exploration in high-dimensional or weakly informative objective landscapes [13, 14, 15].

The generality of the LW foraging hypothesis remains debated. Its optimality depends on assumptions about spatial scale, resource renewal, target density, memory, and boundary effects [16, 17, 18]. In finite landscapes, very long moves can be altered by boundaries or by repeated crossings of already explored regions. In depletable landscapes, local resource consumption alters the effective environment during the search, forcing the forager to balance exploitation of nearby resources against relocation to less-depleted areas. Moreover, Lévy-like trajectories may arise either from intrinsic movement rules or from interactions between the forager and a structured environment, making it difficult to infer the underlying strategy from step-length statistics alone [11, 19]. A closely related alternative is an intermittent search (IS), in which movement alternates between localized search phases and relocation phases [20]. Unlike a strict LW, which is defined by a scale-free move-length law, IS can generate bimodal or multimodal movement statistics, with frequent local displacements and less frequent long transfers. Recent evolutionary simulations that relaxed the assumption of power-law move lengths found that IS can evolve de novo and outcompete LWs under a broad class of foraging conditions [19, 21]. These results suggest that selection may favor a structured mixture of local exploitation and relocation rather than a strict scale-free movement law. Evolutionary models have also supported the idea that Lévy-like search can arise under sparse-resource conditions [21, 22]. Previous work has connected stochastic-search models with classical foraging theory and examined the assumptions and competing interpretations of the Lévy-foraging hypothesis [23, 24]. Within this broader context, the present study considers a specific evolutionary setting with finite, depletable landscapes and compares the resulting trajectories using LW and IS moment models.

Recent random-walk theory also provides tools for studying foraging under depletion. Range-controlled random walks, in which movement rates depend on the number of distinct sites already visited, show how exploration dynamics can change when motion is coupled to accumulated coverage [25, 26, 27]. Starving random-walk models make this connection explicit by asking whether a forager can discover a new food-containing site before a metabolic deadline [28, 29, 30, 31]. In this setting, the new-site discovery interval τ_n, defined as the waiting time between discovery of the n-th and $n+1$-th distinct lattice sites, provides a natural measure of exploration under depletion. These ideas motivate the analysis of waiting-time structure in addition to spatial trajectory shape.

Distinguishing LW and IS dynamics requires more than inspecting apparent heavy tails. Moment-based approaches provide one route to this distinction. In recent work, the second and fourth displacement moments were shown to separate LW and IS dynamics because LWs produce scale-dependent moment growth governed by the Lévy exponent, whereas IS produces a two-timescale structure associated with relocation and local search phases [32, 33]. This provides a quantitative way to test whether evolved trajectories are better described by strict LW dynamics or by IS.

Here, we develop an evolutionary simulation framework that combines genome-encoded movement traits, finite lifetimes, explicit resource depletion, and controlled resource clustering. Agents forage on two-dimensional toroidal lattices containing non-renewable resources distributed either as Lévy dust or uniformly. Selection acts on a fitness function that rewards energetic efficiency and low-redundancy spatial coverage. We analyze how movement genomes evolve across resource environments, classify evolved trajectories using second- and fourth-moment fits to LW and IS models, and quantify waiting-time statistics under depletion. This framework allows us to ask whether adaptive search in finite, depletable landscapes converges toward strict scale-free LW dynamics or toward IS, combining local exploitation with occasional relocation.

2 Methods

2.1 Resource landscapes and depletion

We simulate evolutionary foraging on a discrete two-dimensional torus of size $L \times L$, populated by depletable unit-valued resources. Clustered resource landscapes were generated as Lévy-dust fields, following standard constructions in which successive resource placements are separated by power-law-distributed distances [6, 19, 34]. Resources are non-renewable: once encountered by an agent, the resource at that site is removed for the remainder of that lifetime. The realized trajectory is therefore shaped jointly by the initial spatial resource field and by depletion during search. Resource placement follows a random walk on the torus whose placement-step lengths ℓ are sampled from

$$p_\mu(\ell) = \frac{\ell^{-\mu}}{\sum_{k=1}^{\lfloor L/2 \rfloor} k^{-\mu}}, \qquad \ell = 1, 2, \ldots, \lfloor L/2 \rfloor. \tag{1}$$

Directions are sampled uniformly, and positions are wrapped modulo L, corresponding to periodic boundary conditions. Each reached lattice site is marked as containing a unit resource. The placement walk is run for $A_\rho = \lfloor \rho L^2 \rfloor$ placement attempts. Since the placement walk may return to sites that have already been occupied, repeated hits do not create additional occupied sites. Thus, for Lévy-dust landscapes, ρ specifies the nominal number of placement attempts per lattice site, whereas the realized density is the number of distinct occupied sites divided by L^2. The realized density can therefore differ from the nominal value, especially when repeated placements occur. We consider Lévy-dust exponents $\mu \in \{1.0, 1.5, 2.0, 2.5, 3.0\}$, which control the spatial correlation of resource placement. Smaller μ values allow more long placement steps and produce more spatially dispersed resource fields, whereas larger μ values suppress long placement steps and increase local aggregation. As a reference condition, we also use a uniform environment generated by an independent Bernoulli occupation of lattice sites with probability ρ. In the Lévy-dust fields, repeated placements at the same site leave the occupancy state unchanged. The simulations used in this study consider three linear system sizes, $L \in \{211, 503, 1009\}$. At fixed nominal density ρ, changing μ alters both spatial clustering and realized resource abundance through repeated resource placements. Comparisons across μ therefore reflect changes in both resource geometry and abundance. Table 3 in Appendix B reports the realized occupancies of the saved pre-depletion maps for $L = 503$ at nominal densities $\rho \in \{0.005, 0.010, 0.15, 0.30, 0.45\}$. For each grid size, nominal resource density, and resource environment, the resource field is generated on the torus and used as the initial landscape for the corresponding evolutionary run. Figure 1 shows representative resource maps and best evolved trajectories for the $L = 503$ system, together with the placement-step distributions used to generate the Lévy-dust environments. Additional implementation details with the simulation parameters are provided in Appendix A, Table 2.

2.2 Agent movement and genetic representation

Each agent carries three heritable integer-valued genomes of fixed length G: a step-length genome, a velocity genome, and a turn-angle genome. The admissible values are $s \in \{1, \ldots, \lfloor L/2 \rfloor\}$, $v \in \{1, \ldots, L\}$, $\kappa \in \{0, \ldots, n_\theta - 1\}$, with $n_\theta = 36$, where κ maps to the turning angle $\theta_\kappa = 10^\circ \kappa$. Each genome contains $G = 30$ entries. The genomes are represented as finite lists of admissible trait values rather than as explicit parametric distributions. At each movement decision, the agent samples one entry uniformly from each list. Selection, recombination, and mutation, therefore, act on the empirical frequencies of encoded trait values, allowing the movement statistics to emerge from the evolving genome. At the beginning of a lifetime, the agent is initialized at a resource-bearing lattice site chosen uniformly at random. If no resource-bearing site is available, the initial position is drawn uniformly from the lattice. The initial heading is sampled uniformly. Motion then proceeds on the periodic lattice for the prescribed lifespan T_{life}. At each movement decision, the intended step length and velocity are sampled from the corresponding genomes. With probability $p_{\mathrm{cue}} = 0.5$, the cue branch searches a square window of half-width $r_{\mathrm{cue}} = 60$, clipped at the lattice boundaries. Movement is permitted across the periodic boundary, but sensing does not wrap. If the window contains an undepleted resource, the heading is set directly toward the nearest one by Euclidean distance. Otherwise, a uniformly random heading is selected without falling through to the genome rule. With probability $1 - p_{\mathrm{cue}}$, the inherited turning increment is applied. Thus, p_{cue} is the branch-selection probability, while the frequency of resource-directed turns also depends on resource

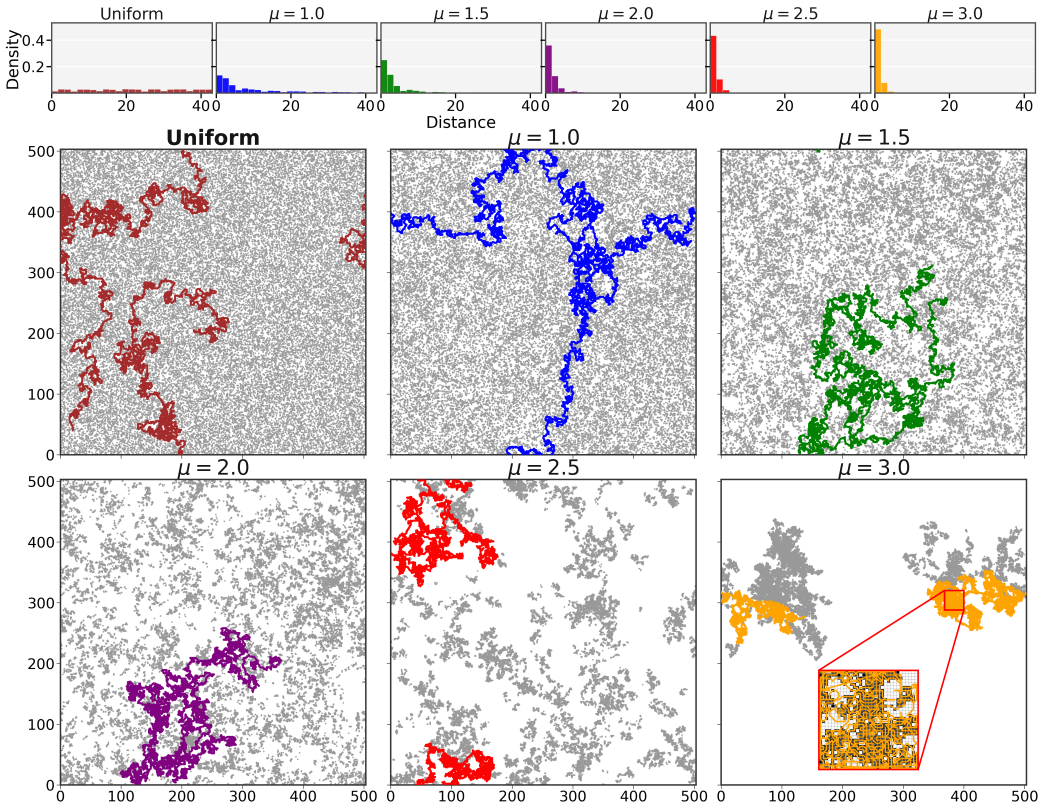


Figure 1: Resource landscapes and best evolved trajectories for the $L = 503$ torus. **Top row**: placement-step distributions used to generate the resource fields, shown for the uniform control and for Lévy-dust environments with $\mu \in \{1.0, 1.5, 2.0, 2.5, 3.0\}$. **Main panels**: resource locations (gray points) overlaid with the trajectory of the best evolved agent in each environment (colored line). Changing μ changes the spatial correlation structure of the resource field by altering the frequency of long placement steps: smaller μ gives heavier-tailed placement steps, whereas larger μ suppresses long placement steps. The inset in the $\mu = 3.0$ panel shows a magnified local patch.

availability. This rule introduces limited local environmental information without giving the agent global knowledge of the resource field [35].

The intended displacement is a straight movement segment of length s, realized on the lattice along the current heading and subject to periodic wrapping. The velocity determines the maximum distance advanced during one movement update within that segment. If s_{rem} denotes the remaining segment length, movement update j advances $\Delta_j = \min(v, s_{\mathrm{rem}})$, followed by $s_{\mathrm{rem}} \leftarrow s_{\mathrm{rem}} - \Delta_j$. An uninterrupted segment therefore requires $\lceil s/v \rceil$ movement updates. The lifespan $T_{\mathrm{life}} = 8000$ counts movement updates, and the accumulated movement cost is $Q(t) = c_{\mathrm{move}} \sum_{j=1}^{t} \Delta_j$, where $c_{\mathrm{move}} = 0.5$. Thus, Δ_j is the scalar advance used to reduce the remaining segment length and to calculate movement cost. It need not equal the Euclidean length of the rounded, pre-wrap coordinate increment. A site counts as visited only when it is the lattice position at the end of a movement update. Sites between consecutive endpoints do not increment $N(t)$ or $R(t)$, and resources on those sites are neither encountered nor consumed. If a resource is present at an update endpoint, the resource is consumed, the cumulative harvested energy $E_{\mathrm{res}}(t)$ is incremented by the prescribed gain, and the current movement segment is terminated at that endpoint. During each lifetime, we record the cumulative movement cost $Q(t)$, the number $N(t)$ of distinct lattice sites visited, and the revisit count $R(t)$. The initial position is included as the first element of the endpoint sequence, and all coverage and waiting-time statistics are calculated from this same sequence. These quantities serve as the state variables in the fitness function defined below.

2.3 Fitness evaluation

Fitness is defined to reward agents that acquire resources at low movement cost while maintaining broad, low-redundancy spatial coverage. We choose the product of an energetic-efficiency term and a coverage-efficiency term as the modelling objective, namely

$$\mathcal{F}(t) = \eta_E(t)\,\eta_C(t). \tag{2}$$

To define the energetic efficiency $\eta_E(t)$, we consider the cumulative movement cost $Q(t)$, calculated from the scalar movement increments Δ_j, and the harvested resource energy $E_{\mathrm{res}}(t)$. This energy is computed as $E_{\mathrm{res}}(t) = \varepsilon n_{\mathrm{res}}(t)$, where ε is the gain from one resource encounter and $n_{\mathrm{res}}(t)$ is the number of resources collected up to time t. The energetic efficiency is $\eta_E(t) = E_{\mathrm{res}}(t)/(E_{\mathrm{res}}(t) + Q(t)) = 1/(1 + Q(t)/E_{\mathrm{res}}(t))$. When $E_{\mathrm{res}}(t) = 0$, we set $\eta_E(t) = 0$, so that agents that collect no resources receive zero energetic efficiency. Note that $0 \le \eta_E(t) \le 1$. This term is large when the agent obtains substantial reward at low locomotion cost and decreases when search becomes energetically inefficient.

To quantify coverage-efficiency, let $N(t)$ be the number of distinct lattice sites visited up to time t, and let $R(t)$ be the number of revisits to sites that have already been visited. Then, similarly to $\eta_E(t)$ we define $\eta_C(t) = N(t)/(N(t) + R(t)) = 1/(1 + R(t)/N(t))$ with $0 \le \eta_C(t) \le 1$. This term approaches one when most visits are to previously unvisited sites and decreases when the trajectory becomes dominated by returns to already explored locations. Combining the two terms gives

$$\mathcal{F}(t) = \frac{1}{1 + \frac{Q(t)}{E_{\mathrm{res}}(t)}} \cdot \frac{1}{1 + \frac{R(t)}{N(t)}}. \tag{3}$$

By construction, $0 \le \mathcal{F}(t) \le 1$. High fitness, therefore, requires both profitable energy acquisition and efficient spatial coverage. In the simulations, the quantities entering Eq. (3) are tracked directly during each trajectory. The harvested energy $E_{\mathrm{res}}(t)$ is updated whenever a resource is encountered and consumed. The cumulative cost $Q(t)$ is incremented according to the realized motion under the fixed movement-cost rule. Distinct visited sites are recorded explicitly, and each return to an already visited site increments the revisit count $R(t)$. Fitness is evaluated at the end of the lifetime and used for fitness-proportionate reproduction with elitism, recombination, and bounded Gaussian mutation.

2.4 Classifying evolved trajectories: intermittent or Lévy?

To determine whether the evolved trajectories are better described by an IS model or by an LW model, we fitted theoretical displacement moments to empirical moments computed from the simulated paths. The procedure follows the log-moment classification framework introduced in Ref. [32], with fitting and scoring routines implemented in the IntLevPy analysis library [33]. For each trajectory, we computed displacements over lag t_s, $\Delta\mathbf{r}(t; t_s) = \mathbf{r}(t + t_s) - \mathbf{r}(t)$, using unwrapped coordinates to avoid artificial shortening at periodic boundary crossings. The empirical radial moments were

$$m_k^{\mathrm{obs}}(t_s) = \left\langle \|\Delta\mathbf{r}(t; t_s)\|^k \right\rangle_t, \qquad k \in \{2, 4\}, \tag{4}$$

where the average is taken over all admissible time origins t along the trajectory. Fitting used a set of $n_{\mathrm{lag}} = 28$ unique lag times t_s, generated as $t_s = \lfloor 2^{n/10} \rfloor$ for $n = 0, \ldots, 53$, with all lag times measured in movement updates.

The LW model is defined by straight flights at constant speed, with flight durations drawn from a heavy-tailed distribution,

$$\psi(\tau_{\mathrm{f}}) = \frac{\gamma}{\tau_0} \left(1 + \frac{\tau_{\mathrm{f}}}{\tau_0}\right)^{-(1+\gamma)}, \tag{5}$$

parameterized by (v_L, γ, τ_0). The IS model alternates between a diffusive phase of diffusivity D and a ballistic phase of speed v_B, with Poisson switching rates λ_{BD} and λ_{DB}. For the IS model, explicit analytical expressions for the second and fourth displacement moments are available [32]. See Appendix C. For the LW model, the corresponding moment behavior in log-space was evaluated within the same fitting framework used in Ref. [32] and in the associated implementation software library [33]. For each model, parameters were estimated by minimizing the mean squared discrepancy in logarithmic moment space,

$$\mathcal{J} = \frac{1}{2n_{\mathrm{lag}}} \sum_{t_s} \sum_{k \in \{2,4\}} \left[\log m_k^{\mathrm{obs}}(t_s) - \log m_k^{\mathrm{mod}}(t_s)\right]^2, \tag{6}$$

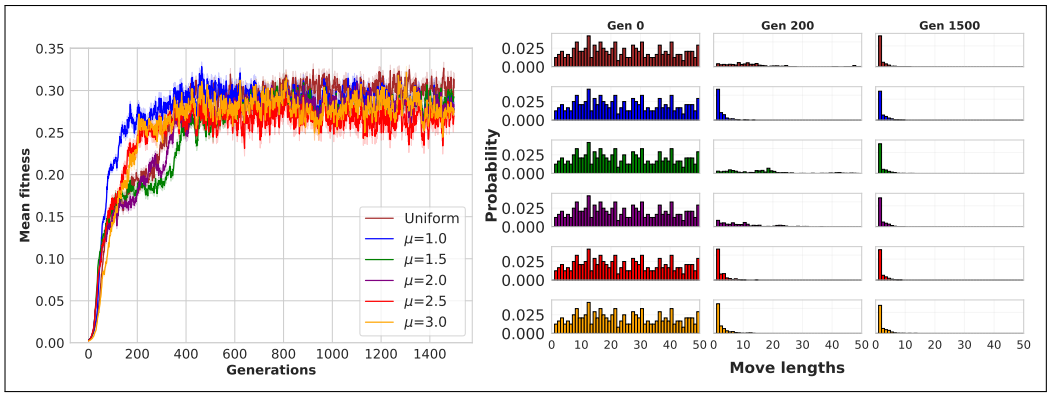


Figure 2: Evolution of fitness and the step lengths on the 503×503 grid. **Left:** mean fitness over 1500 generations for each resource environment, with shaded standard errors across agents within a single evolutionary run at each generation. **Right:** distributions of step-length genome entries at generations 0, 200, and 1500. Selection rapidly concentrates the encoded step-length pool at short values while retaining a sparse tail of longer entries.

subject to the corresponding physical parameter constraints. Working in log-moment space reduces the dominance of large-amplitude moment values and gives a more balanced comparison across lag scales. Goodness of fit was quantified using the adjusted coefficient of determination computed in logarithmic moment space, accounting for the three free parameters of LW and four of IS. For each trajectory and model, the adjusted R^2 values for m_2 and m_4 were averaged to give $\bar{R}^2_{\mathrm{IS}}$ and $\bar{R}^2_{\mathrm{LW}}$. We then defined the model-comparison score

$$\Gamma = \bar{R}^2_{\mathrm{IS}} - \bar{R}^2_{\mathrm{LW}}. \tag{7}$$

Thus, $\Gamma > 0$ indicates that the IS model provides a better moment-level description of the trajectory, whereas $\Gamma < 0$ favors the LW model. The classification is therefore a comparison between two idealized dynamical descriptions, not evidence that agents explicitly implement either model; summary statistics are computed over stored generation-wise best-agent trajectories. The classification concerns realized trajectories shaped jointly by the genomes, sensory cue, depletion, and encounter-triggered stopping. Synthetic-trajectory classification and IS parameter recovery are reported in Appendix B.

3 Results: evolved search dynamics

We analyzed evolved trajectories and movement traits after 1500 generations of selection across all resource environments. Across conditions, fitness increased early in evolution, step-length genomes shifted toward short displacements, and the evolved trajectories were better approximated by the IS model than by the strict LW model. We first describe the spatial trajectories and evolutionary changes in the genomes, and then compare the evolved trajectories with the IS and LW models using the second and fourth-order displacement moments. The trajectories in Figure 1 show the best evolved individual on the 503×503 torus, overlaid on the corresponding resource maps with $\rho = 0.15$. The overlays show spatially heterogeneous movement, with dense local path segments and occasional larger displacements visible in several environments. We treat this as a qualitative description of the trajectories; the quantitative distinction between LW and IS dynamics is made using second- and fourth-order displacement moments. As the resource field changes with μ, the spatial organization of the trajectories changes as well, reflecting the interaction between the evolved movement genome, resource depletion, and movement termination upon resource encounter. Because resource encounters can truncate planned movement segments, realized displacement statistics need not coincide with the raw genome-encoded step-length lists.

Figure 2 summarizes the evolutionary dynamics on the 503×503 grid. Mean fitness rises rapidly during the early generations and then fluctuates around an environment-dependent plateau. The curves do not show a simple monotonic ordering with μ, indicating that final performance reflects the combined effects of resource geometry, depletion, local cueing, and movement cost. The genome histograms show a clearer and more robust trend: initially broad step-length lists become concentrated

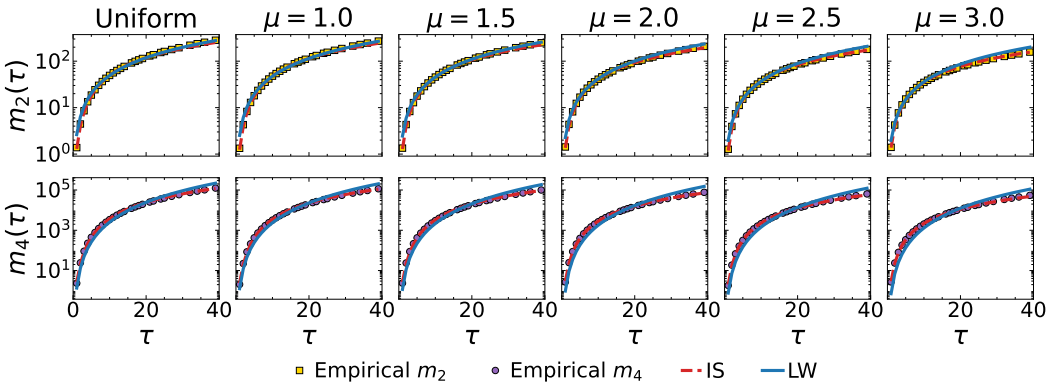


Figure 3: Moment-based comparison of evolved trajectories on the 503×503 grid. Empirical second moments m_2 and fourth moments m_4 are compared with best-fit IS and LW models. Columns correspond to the six resource environments. Across environments, IS gives a closer fit to the empirical moment curves.

at short steps by generation 200 and remain strongly short-biased at generation 1500, while retaining a sparse tail of longer entries. These histograms describe the encoded sampling pool from which step lengths are drawn, not the realized displacement distribution along the path. A complementary stabilization analysis is provided in Appendix A. There, the best-fitness trajectories are shown together with the Jensen–Shannon divergence $\mathrm{JS}(P_g(s), P_0(s))$ of the population-level step-length distribution from its initial state. Both quantities show rapid early change followed by an operational plateau, with plateau generations estimated between $g = 390$ and $g = 547$ on the 503×503 grid using a rolling-slope criterion.

Figure 3 compares empirical second and fourth displacement moments with the best-fit IS and LW models for the 503×503 grid ($\rho = 0.15$). The two models often appear close at the level of the second moment, but differences are clearer in the fourth moment, which is more sensitive to rare large displacements. Across all environments, the IS model provides a closer moment-level description. This visual pattern is consistent with the quantitative model-comparison scores in Table 1 for every environment and grid size, $\Gamma = \bar{R}^2_{\mathrm{IS}} - \bar{R}^2_{\mathrm{LW}} > 0$. Thus, within this moment-based classification framework, the evolved trajectories are better approximated by the IS than by the strict LW model. At $L = 503$ and $\rho = 0.15$, five independent evolutionary runs per environment reproduced the preference for IS, with $\Gamma > 0$ in all 30 run–environment combinations. Positive Γ was also observed across all six environments in the single-run density sweep at $\rho \in \{0.005, 0.010, 0.15, 0.30, 0.45\}$ (Appendix B, Fig. 8).

Re-evolution under an ablated objective $F = \eta_E$ retained 94.7%–100.1% of the baseline coverage efficiency and yielded $\Gamma > 0$ in all six environments (Appendix Table 4). Setting $p_{\mathrm{cue}} = 0$ during fixed-genome evaluation also preserved $\Gamma > 0$, although the margin decreased in the uniform and $\mu \in \{1.0, 1.5, 2.0\}$ environments and increased for $\mu \in \{2.5, 3.0\}$ (Appendix Table 5). Because the genomes were not re-evolved without the cue, the latter analysis measures its effect on expressed trajectories rather than on evolutionary outcomes.

Figure 4 compares movement traits and return-interval statistics of the best evolved agents across the three grid sizes. The top row shows velocity-genome entries. The median velocities range from 1.0 to 9.5 lattice units per update across the displayed conditions, indicating that the characteristic speed scale is not fixed universally, but varies with resource environment and system size. The middle row shows the distribution of turn-angle increments recorded along the best evolved trajectories. Angles are represented modulo 360° and binned in 10° intervals. Probability mass near 0° or 360° corresponds to nearly forward motion, whereas mass at intermediate angles corresponds to reorientation events. The bottom row shows the running maximum of return intervals, $M_j^{\mathrm{return}} = \max\{\Delta t_1^{\mathrm{return}}, \ldots, \Delta t_j^{\mathrm{return}}\}$, as a function of the return-interval index j. Here $\Delta t_j^{\mathrm{return}}$ denotes the j-th observed return interval between successive visits to the same lattice site. The staircase form is expected for a record process: M_j^{return} changes only when a newly observed return interval exceeds

Table 1: Model-comparison summary across environments and grid sizes at $\rho = 0.15$. For each case, we report $\Gamma = \bar{R}^2_{\mathrm{IS}} - \bar{R}^2_{\mathrm{LW}}$, together with the adjusted coefficients of determination. For $L = 211$ and $L = 1009$, values are mean $\pm$ standard deviation over the final 500 best-agent trajectories from one evolutionary run. For $L = 503$, values are mean $\pm$ standard deviation across five independent run-level means, each computed from the final 500 best-agent trajectories. The runs used independent resource maps and GA initializations. Positive Γ indicates preference for the IS model. All results are reported with significant figures in accordance with their respective standard deviations.

Grid	Environment	Γ	$\bar{R}^2_{\mathrm{IS}}$	$\bar{R}^2_{\mathrm{LW}}$
211×211	uniform	0.028 ± 0.004	0.9985 ± 0.0004	0.970 ± 0.004
211×211	$\mu = 1.0$	0.019 ± 0.002	0.9984 ± 0.0002	0.980 ± 0.002
211×211	$\mu = 1.5$	0.022 ± 0.003	0.9981 ± 0.0003	0.976 ± 0.003
211×211	$\mu = 2.0$	0.024 ± 0.005	0.9993 ± 0.0002	0.975 ± 0.005
211×211	$\mu = 2.5$	0.030 ± 0.007	0.9987 ± 0.0005	0.969 ± 0.007
211×211	$\mu = 3.0$	$\mathbf{0.08 \pm 0.02}$	0.9986 ± 0.0006	0.92 ± 0.02
503×503	uniform	0.017 ± 0.002	0.9983 ± 0.0004	0.982 ± 0.002
503×503	$\mu = 1.0$	0.015 ± 0.005	0.9986 ± 0.0004	0.984 ± 0.005
503×503	$\mu = 1.5$	0.018 ± 0.006	0.9987 ± 0.0004	0.981 ± 0.006
503×503	$\mu = 2.0$	0.030 ± 0.002	0.9991 ± 0.0002	0.969 ± 0.002
503×503	$\mu = 2.5$	0.04 ± 0.01	0.9992 ± 0.0002	0.960 ± 0.009
503×503	$\mu = 3.0$	$\mathbf{0.08 \pm 0.03}$	0.9991 ± 0.0002	0.92 ± 0.03
1009×1009	uniform	0.017 ± 0.001	0.9983 ± 0.0003	0.981 ± 0.001
1009×1009	$\mu = 1.0$	0.019 ± 0.001	0.9985 ± 0.0002	0.979 ± 0.001
1009×1009	$\mu = 1.5$	0.022 ± 0.006	0.9987 ± 0.0006	0.977 ± 0.006
1009×1009	$\mu = 2.0$	0.031 ± 0.005	0.9992 ± 0.0003	0.968 ± 0.004
1009×1009	$\mu = 2.5$	0.06 ± 0.02	0.9993 ± 0.0002	0.94 ± 0.02
1009×1009	$\mu = 3.0$	$\mathbf{0.11 \pm 0.04}$	0.9990 ± 0.0006	0.89 ± 0.04

all previous return intervals. The corresponding upper-tail ratios $q_{0.95}/\mathrm{median}$, reported in Appendix Table 6, confirm that return intervals have broad upper tails across the tested conditions. These curves quantify recurrence to previously visited sites and are distinct from the new-site discovery intervals τ_n, which measure the waiting time to discover previously unvisited sites.

Taken together, the short-biased step-length genomes, structured turning statistics, broad return-interval distributions, and consistently positive model-comparison scores support a common interpretation. In finite, depleting landscapes, the evolved search dynamics are better described as IS, combining local exploitation with occasional relocation, than as strict scale-free LW dynamics. Additional waiting-time analyses in Appendix B show that return and resource-hit intervals have broader upper tails than new-site discovery intervals, further supporting this interpretation.

4 Discussion and conclusions

We studied how search strategies evolve in finite, depletable resource landscapes when movement rules are inherited but not prescribed as a fixed parametric distribution. Across all grid sizes and resource environments, evolution shifted the step-length genomes from broadly distributed initial trait lists toward short-step-biased distributions with a sparse tail of longer entries. At the trajectory level, the resulting paths are consistent with local scanning interrupted by occasional relocation. Moment-based model comparison consistently favored IS over LW dynamics, with positive values of Γ across all environments and system sizes. Thus, within the present framework, adaptive search in finite depletion-driven landscapes is better described by IS than by strict scale-free LWs. The single-run sparse-density tests retained the IS fit preference. In a separate single-run ablation, agents were re-evolved using $F = \eta_E$, and Γ remained positive in all six environments. In the fixed-genome cue-off evaluation, the mean score difference remained positive, but its magnitude changed in opposite directions across environments. These findings show that the coverage term is not solely responsible for the classification and that sensing affects the expressed trajectories. They do not establish cue-free evolution or isolate the effects of depletion and encounter-triggered stopping.

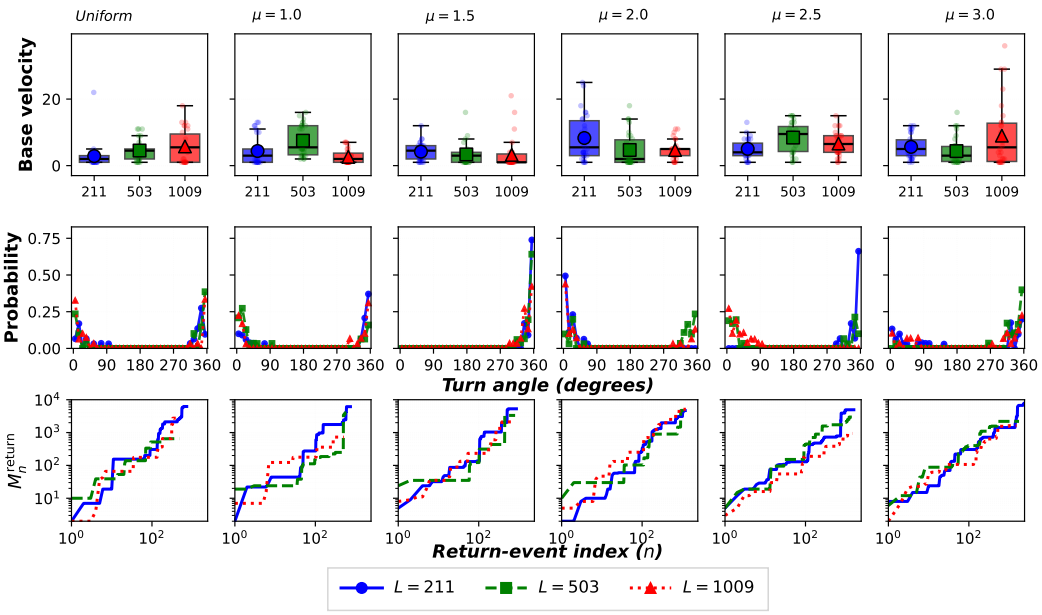


Figure 4: Movement traits and return-interval statistics of the best evolved agents across grid sizes after 1500 generations. Columns correspond to resource environments. **Top row:** velocity-genome entries of the best evolved agents for $L = 211$, 503, and 1009. **Middle row:** turn-angle increments recorded along the best evolved trajectories. **Bottom row:** running maximum of return intervals M_j^{return} versus return-event index j.

This result is consistent with the view that Lévy-like search is not a universal optimum, but depends on ecological and physical constraints such as resource renewal, target density, memory, boundary effects, landscape heterogeneity, and sampling scale [5, 7, 16, 18, 36, 37, 38]. In the present model, resources are consumed upon encounter and do not regenerate during a lifetime. Under such conditions, a forager that has just encountered a resource may benefit from continued local search, because nearby resources are more likely to occur in clustered landscapes. However, local exploitation eventually becomes inefficient as the neighborhood is depleted. Occasional relocations are potentially useful because they move the agent away from over-explored regions. This interpretation is also consistent with movement-ecology work showing that animals in heterogeneous landscapes often alternate between intensive movement in profitable areas and more extensive movement between them [39]. The evolved strategy reflects this balance: short movements reduce cost and support local exploitation, while rarer longer movements reduce stagnation. The comparison with IS is also consistent with recent evolutionary work showing that, when the move-length distribution is not constrained to be a power law, evolved search can favor IS rather than strict Lévy flight [19]. Our results extend this idea to explicitly depletable two-dimensional resource landscapes with genome-encoded step length, velocity, and turning traits. The evolved genomes do not produce a pure scale-free law. Instead, selection concentrates the encoded step-length pool at short values while preserving a small relocation component. This is closer to the structure expected for IS, where localized scanning and relocation are distinct functional components [20, 40], than to a strict Lévy walk defined by a single heavy-tailed flight distribution.

The displacement-moment analysis provides a quantitative way to distinguish these alternatives. This is important because step-length or path-shape inspection alone can be sensitive to sampling choices and environmental structure, and can therefore blur the distinction between intrinsic movement rules and environment-induced movement patterns [37, 38, 39]. The second and fourth moments probe different aspects of the trajectory: the second moment captures the typical spread of displacement with lag, whereas the fourth moment is more sensitive to rare large displacements. Across environments, the IS model gives a closer moment-level description, and the preference remains positive across the three grid sizes. This supports the preference for the IS description in this two-model comparison,

without identifying a unique generating mechanism or establishing superior foraging performance. The result is also consistent with the moment-based distinction between Lévy and IS developed in Refs. [32, 33]. The waiting-time statistics provide a complementary view of the same dynamics. The running maximum of return intervals shows that extreme returns to already visited sites arise as rare record events, while the appendix analysis separates recurrence, resource encounter timing, and new-site discovery. Return intervals and resource-hit intervals have broader upper tails than new-site discovery intervals, indicating that the strongest temporal intermittency appears in recurrence and resource-encounter timing rather than in long pauses in spatial discovery. This distinction is important because the visitation-statistics quantity τ_n refers to the time needed to discover a new site after n distinct sites have already been visited [31]. In our analysis, τ_n is therefore treated separately from return intervals to previously visited sites.

These findings connect evolutionary foraging with random-walk models of depletion-controlled search. Starving random walks and related visitation-statistics models emphasize that survival depends not only on the total number of sites visited, but also on the temporal structure of new-site discovery [29, 30, 31]. Our simulations do not implement starvation through a hard metabolic deadline, but they share the same depletion principle: continued success requires movement into regions where resources remain available. The evolved strategies therefore occupy an intermediate position between idealized random-walk theory and adaptive foraging. They are shaped by local resource geometry, depletion, movement cost, and heritable variation in movement traits. The results also suggest a possible design principle for resource-limited search and optimization. Movement rules dominated by short displacements but punctuated by occasional relocations may be useful in heterogeneous objective landscapes, especially when local improvement is likely but prolonged exploitation becomes inefficient [7, 13, 14]. In autonomous search, a comparable rule would combine local scanning after detection with occasional relocation to avoid spending excessive time in depleted or low-yield regions. The reactive stopping rule in our simulations, in which movement segments terminate upon resource encounter, is a simple instance of this principle and may be relevant to energy-limited robotic or algorithmic exploration [6].

Several features of the model limit the interpretation of these results. The search domain is finite and periodic, so the simulations do not access the asymptotic regimes assumed in many analytical LW or random-walk results. Resource landscapes are fixed during a lifetime except for depletion, and no resource renewal is included. The evolutionary protocol also uses fixed population size, mutation parameters, and genome length across all environments. Independent replication is limited to $L = 503$ and $\rho = 0.15$. The other grid sizes and density conditions remain single-run tests. Finally, the simulations are single-agent evaluations and do not include social information, competition, or collective effects, which can alter the balance between exploration and exploitation [12, 41, 42]. Future work should examine larger domains, explicit resource renewal, and alternative sensing rules to determine when intermittent strategies persist and when scale-free movement becomes advantageous. Multi-agent extensions would also be useful for studying how cooperation, interference, or social learning modifies the evolved search strategy. More broadly, although the present work is a controlled simulation study rather than a deployable system, search rules of this kind could support energy-aware autonomous exploration and ecological monitoring. Still, they would require privacy, safety, and oversight safeguards if used in sensitive human or environmental settings. In summary, evolution on finite, depletable landscapes drives movement genomes toward short-step-biased trait distributions with occasional relocation, and the IS dynamics more accurately describe the resulting trajectories than strict LW within the moment-based classification used here. At $L = 503$ and $\rho = 0.15$, independent runs support a reproducible IS–LW fit preference, conditional on the model's sensing, depletion, and encounter-stopping rules.

Acknowledgments

We thank Pedro Lencastre for providing synthetic trajectories used to validate our moment-based classifier. The authors acknowledge funding from the Research Council of Norway under grant number 335940 for the project *Virtual-Eye*. The work of AY was also supported by the Center of Excellence, *Integreat*, project number 332645. PGL acknowledges support by *MishMash* Centre for AI and Creativity, Project Number 357438.

References

[1] Eric L. Charnov. Optimal foraging, the marginal value theorem. *Theoretical Population Biology*, 9(2):129–136, 1976.

[2] Dean Mobbs, Pete C. Trimmer, Daniel T. Blumstein, and Peter Dayan. Foraging for foundations in decision neuroscience: insights from ethology. *Nature Reviews Neuroscience*, 19(7):419–427, July 2018.

[3] L. Wang, T. Wang, X.-W. Zhang, X.-F. Lin, J. Li, J.-B. Liao, et al. Evolution of foraging behaviour induces variable complexity-stability relationships in mutualist-exploiter-predator communities. *PLoS Computational Biology*, 21(7):e1013245, 2025.

[4] Nicolas Chaumont and Christoph Adami. Evolution of sustained foraging in three-dimensional environments with physics. *Genetic Programming and Evolvable Machines*, 17(4):359–390, 2016.

[5] G. M. Viswanathan, S. V. Buldyrev, S. Havlin, M. G. E. da Luz, E. P. Raposo, and H. E. Stanley. Optimizing the success of random searches. *Nature*, 401(6756):911–914, 1999.

[6] Gandhimohan M. Viswanathan, Marcos G. E. da Luz, Ernesto P. Raposo, and H. Eugene Stanley. *The physics of foraging: an introduction to random searches and biological encounters*. Cambridge University Press, 2011.

[7] V. Zaburdaev, S. Denisov, and J. Klafter. Lévy walks. *Rev. Mod. Phys.*, 87:483–530, Jun 2015.

[8] Brieuc Guinard and Amos Korman. Intermittent inverse-square lévy walks are optimal for finding targets of all sizes. *Science Advances*, 7(15):eabe8211, 2021.

[9] Malcolm D. O'Toole, Mary-Anne Lea, Christophe Guinet, Robert Schick, and Mark A. Hindell. Foraging strategy switch of a top marine predator according to seasonal resource differences. *Frontiers in Marine Science*, Volume 2 - 2015, 2015.

[10] Ludovica Luisa Vissat, Shlomo Cain, Sivan Toledo, Orr Spiegel, and Wayne M. Getz. Categorizing the geometry of animal diel movement patterns with examples from high-resolution barn owl tracking. *Movement Ecology*, 11(1):15, Mar 2023. Received: 20 May 2022; Accepted: 21 January 2023; Published: 21 March 2023.

[11] A. O. Gautestad and A. Mysterud. The lévy flight foraging hypothesis: forgetting about memory may lead to false verification of brownian motion. *Movement Ecology*, 1:9, 2013.

[12] Ketika Garg, Paul E Smaldino, and Christopher T Kello. Evolution of explorative and exploitative search strategies in collective foraging. *Collective Intelligence*, 3(1):26339137241228858, 2024.

[13] Vincenzo Fioriti, Fabio Fratichini, Stefano Chiesa, and Claudio Moriconi. Levy foraging in a dynamic environment – extending the levy search. *International Journal of Advanced Robotic Systems*, 12(7), 07 2015.

[14] Dario A. Leon and Augusto Gonzalez. Mutations as lévy flights. *Scientific Reports*, 11:9889, 2021.

[15] Quanqin He, Hao Liu, Guiyan Ding, and Liangping Tu. A modified lévy flight distribution for solving high-dimensional numerical optimization problems. *Mathematics and Computers in Simulation*, 204:376–400, 2023.

[16] Kun Zhao, Raja Jurdak, Jiajun Liu, David Westcott, Branislav Kusy, Hazel Parry, Philipp Sommer, and Adam McKeown. Optimal lévy-flight foraging in a finite landscape. *Journal of The Royal Society Interface*, 12(104):20141158, 03 2015.

[17] A.M. Reynolds. Cooperative random lévy flight searches and the flight patterns of honeybees. *Physics Letters A*, 354(5):384–388, 2006.

[18] Justin C. Tzou and Leo Tzou. Counterexample to the lévy flight foraging hypothesis in the narrow capture framework. *Phys. Rev. Res.*, 6:023274, Jun 2024.

[19] Winston Campeau, Andrew M. Simons, and Brett Stevens. Intermittent search, not strict lévy flight, evolves under relaxed foraging distribution constraints. *The American Naturalist*, 203(4):513–527, 2024.

[20] O. Bénichou, C. Loverdo, M. Moreau, and R. Voituriez. Intermittent search strategies. *Rev. Mod. Phys.*, 83:81–129, Mar 2011.

[21] Winston Campeau, Andrew M. Simons, and Brett Stevens. The evolutionary maintenance of lévy flight foraging. *PLOS Computational Biology*, 18(1):1–26, 01 2022.

[22] Marina E. Wosniack, Marcos C. Santos, Ernesto P. Raposo, Gandhi M. Viswanathan, and Marcos G. E. da Luz. The evolutionary origins of lévy walk foraging. *PLOS Computational Biology*, 13(10):1–31, 10 2017.

[23] F Bartumeus and J Catalan. Optimal search behavior and classic foraging theory. *Journal of Physics A: Mathematical and Theoretical*, 42(43):434002, oct 2009.

[24] Rainer Klages. *Search for Food of Birds, Fish, and Insects*, pages 53–74. Springer International Publishing, Cham, 2023.

[25] Sidney Redner. *A Guide to First-Passage Processes*. Cambridge University Press, Cambridge, UK, 2001.

[26] Elliott W. Montroll and George H. Weiss. Random walks on lattices. ii. *Journal of Mathematical Physics*, 6(2):167–181, 02 1965.

[27] L. Régnier, O. Bénichou, and P. L. Krapivsky. Range-controlled random walks. *Phys. Rev. Lett.*, 130:227101, May 2023.

[28] O. Bénichou and S. Redner. Depletion-controlled starvation of a diffusing forager. *Phys. Rev. Lett.*, 113:238101, Dec 2014.

[29] Daniel ben Avraham and Shlomo Havlin. *Diffusion and Reactions in Fractals and Disordered Systems*. Cambridge University Press, 2000.

[30] Shlomo Havlin and Daniel Ben-Avraham. Diffusion in disordered media. *Advances in Physics*, 36(6):695–798, 1987.

[31] Léo Régnier, Maxim Dolgushev, and Olivier Bénichou. From maximum of inter-visit times to starving random walks. *Phys. Rev. Lett.*, 132:127101, Mar 2024.

[32] Pedro Lencastre, Yurii S. Bystryk, Anis Yazidi, Sergey Denisov, and Pedro G. Lind. Dynamical law behind eye movements: Distinguishing between lévy and intermittent strategies. *Phys. Rev. Res.*, 7:043042, Oct 2025.

[33] Shailendra Bhandari, Pedro Lencastre, Sergey Denisov, Yurii S. Bystryk, and Pedro G. Lind. IntLevPy: A python library to classify and model intermittent and lévy processes. *SoftwareX*, 31:102334, 2025.

[34] P.A. Burrough. Fractal dimensions of landscapes and other environmental data. *Nature*, 294:240–242, 1981.

[35] Ben C. Nolting, Travis M. Hinkelman, Chad E. Brassil, and Brigitte Tenhumberg. Composite random search strategies based on non-directional sensory cues. *Ecological Complexity*, 22:126–138, 2015.

[36] Nicolas E. Humphries, Henri Weimerskirch, Nuno Queiroz, Emily J. Southall, and David W. Sims. Foraging success of biological lévy flights recorded in situ. *Proceedings of the National Academy of Sciences*, 109(19):7169–7174, 2012.

[37] Edward A Codling, Michael J Plank, and Simon Benhamou. Random walk models in biology. *Journal of The Royal Society Interface*, 5(25):813–834, 04 2008.

[38] Michael J. Plank and Edward A. Codling. Sampling rate and misidentification of lévy and non-lévy movement paths. *Ecology*, 90(12):3546–3553, 2009.

[39] F. Barraquand and S. Benhamou. Animal movements in heterogeneous landscapes: identifying profitable places and homogeneous movement bouts. *Ecology*, 89(12):3336–3348, 2008.

[40] Donald L. Kramer and Robert L. McLaughlin. The behavioral ecology of intermittent locomotion. *American Zoologist*, 41(2):137–153, 04 2001.

[41] Nils Kolling and Thomas Akam. (reinforcement?) learning to forage optimally. *Current Opinion in Neurobiology*, 46:162–169, 2017. Computational Neuroscience.

[42] Ketika Garg, Christopher T. Kello, and Paul E. Smaldino. Individual exploration and selective social learning: balancing exploration–exploitation trade-offs in collective foraging. *Journal of The Royal Society Interface*, 19(189):20210915, 04 2022.

[43] M.L. Menéndez, J.A. Pardo, L. Pardo, and M.C. Pardo. The jensen-shannon divergence. *Journal of the Franklin Institute*, 334(2):307–318, 1997.

A The evolutionary algorithm: specifications

We use a generational genetic algorithm to evolve a population of N_{pop} haploid agents over n_{gen} generations. Each agent carries three independent integer-valued genomes of fixed length G, corresponding to step length, velocity, and turning-angle preference:

$$\begin{aligned}
&\mathbf{g}^{(s)} \in \{1, 2, \ldots, \lfloor L/2 \rfloor\}^G, && \text{step-length genome,}\\
&\mathbf{g}^{(v)} \in \{1, 2, \ldots, v_{\max}\}^G, && \text{velocity genome,}\\
&\mathbf{g}^{(\theta)} \in \{0, 1, \ldots, n_\theta - 1\}^G, && \text{turn-angle genome.}
\end{aligned}$$

Here, L is the linear grid size, $v_{\max} = L$ is the maximum admissible velocity, and $n_\theta = 36$ is the number of turning-angle bins, corresponding to angular increments $\theta_\kappa = 10^\circ\kappa$, $\kappa = 0, 1, \ldots, 35$. The genomes are represented as finite lists of admissible trait values rather than as explicit probability distributions. During trajectory generation, phenotypic decisions are obtained by sampling uniformly from these lists. Evolution, therefore, acts on the empirical frequencies of encoded trait values, allowing the movement statistics to emerge from selection, recombination, and mutation.

Each lifetime is simulated using a fresh copy of the initial pre-depletion resource map associated with that evolutionary run, so resource depletion is reset between agent evaluations. At the beginning of each lifetime, an agent is initialized at a resource-bearing lattice site chosen uniformly at random. If no resource-bearing site is available, the initial position is drawn uniformly from the lattice. Heading selection follows the complete rule in Sec. 2.2. With probability (p_{cue}), the cue branch searches the clipped square sensing window. If the window contains an undepleted resource, the heading is set directly toward the nearest resource by Euclidean distance. If the window contains no undepleted resource, a uniformly random heading is selected without returning to the genome branch. With probability $(1 - p_{\mathrm{cue}})$, the genome-sampled turning increment is applied. Only the lattice position at the end of each movement update is checked for resources. If an undepleted resource is present at that endpoint, it is consumed, the harvested energy is incremented, and the current movement segment terminates at that endpoint. Table 2 lists the set of simulation and evolutionary parameters used in the reported experiments.

Fitness evaluation

Fitness is evaluated using the definition given in Sec. 2.3. For each agent, the quantities entering Eq. (3) are accumulated along the realized trajectory, and the terminal value of $\mathcal{F}$ is used as the reproductive score. Harvested energy is updated whenever a resource is encountered and consumed. Movement cost is accumulated from the scalar movement increment (Δ_j). Distinct lattice sites are tracked explicitly, and each return to an already visited site increments the revisit count. Thus, fitness is assigned at the end of the lifetime of each individual.

Figure 5 summarizes three aspects of the evolutionary dynamics on the 503×503 grid. The left panel shows that the best fitness rises rapidly during the early generations and then fluctuates around an environment-dependent plateau. Because mutation and recombination remain active throughout

Table 2: Simulation and evolutionary parameters used for the main experiments. All main results use resource density $\rho = 0.15$. The additional 503×503 density robustness analysis uses $\rho \in \{0.005, 0.010, 0.30, 0.45\}$

.

Parameter	Value
Grid sizes L	$211,\ 503,\ 1009$
Environment types	uniform, $\mu = 1.0,\ 1.5,\ 2.0,\ 2.5,\ 3.0$
Resource density ρ	0.15 for all grid sizes, with $0.005,\ 0.010,\ 0.30,\ 0.45$ additionally tested at $L = 503$
Resource value	1.0
Genome length G	30
Population size N_{pop}	40
Number of generations n_{gen}	1500
Step-length range	$1, \ldots, \lfloor L/2 \rfloor$
Velocity range	$1, \ldots, L$
Turn-angle bins	$0^\circ, 10^\circ, \ldots, 350^\circ$
Sensory cue	$p_{\mathrm{cue}} = 0.5,\ r_{\mathrm{cue}} = 60$, using a square window clipped
Lifespan T_{life}	8000
Movement cost coefficient c_{move}	0.5
Mutation count per genome n_{mut}	6 loci
Number of angle bins n_θ	36
Resource gain ε	1.0
Elitism fraction e_{elite}	0.1
Elite count N_{elite}	$\lfloor e_{\mathrm{elite}} N_{\mathrm{pop}} \rfloor$
Mutation scale σ_{mut}	2
Randomization	seed 42 for the single-run conditions, with independent resource maps and GA initializations for the five replicated $L = 503, \rho = 0.15$ runs
Movement boundary condition	periodic

evolution, this plateau should be interpreted as an operational stabilization regime rather than as exact convergence.

The middle panel shows the cumulative Jensen–Shannon divergence [43] $\mathrm{JS}(P_g(s), P_0(s))$, where $P_g(s)$ is the normalized population-level distribution of genome-encoded step lengths at generation g. The divergence rises rapidly and then saturates, indicating that selection reorganizes the step-length genome early in evolution and subsequently maintains a comparatively stable evolved distribution. To quantify stabilization, we used a rolling-slope criterion. For each environment, rolling linear slopes were computed over a window of 100 generations for both best fitness and $\mathrm{JS}(P_g(s), P_0(s))$. The plateau generation was defined as the first generation after $g = 150$ for which both absolute slopes remained below fixed thresholds for 50 consecutive generations. With thresholds 2×10^{-4} for fitness and 5×10^{-4} for JS divergence, the estimated plateau generations ranged from $g = 390$ to $g = 547$ across environments. These values support the interpretation of rapid early adaptation followed by stabilization, without implying exact convergence.

The right panel decomposes the final best-agent fitness into its energetic and coverage components, η_E and η_C. Since $\mathcal{F} = \eta_E \eta_C$, these values show how final performance is assembled from energy acquisition relative to movement cost and from low-redundancy spatial coverage. Across environments, η_C remains high but decreases from 0.94 in the uniform case to 0.78 at $\mu = 3.0$, whereas η_E ranges from 0.42 to 0.51. Thus, final fitness reflects both energetic cost and the increasing burden of revisits in some environments.

Selection, recombination, and mutation

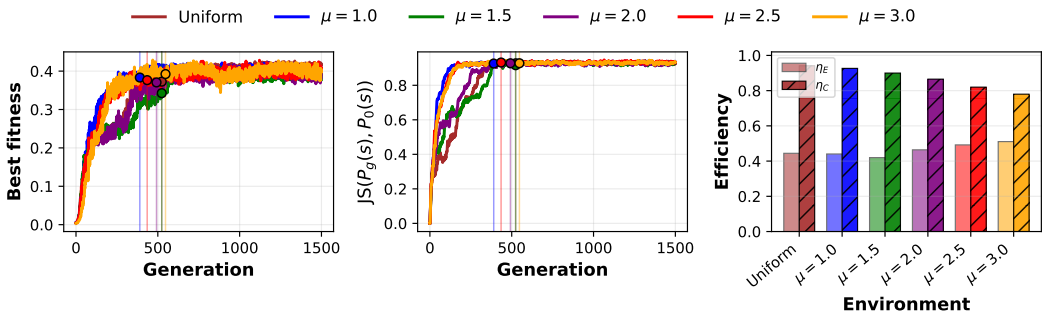


Figure 5: Evolutionary stabilization on the 503×503 grid. Left: best fitness as a function of generation. Middle: JS divergence $\mathrm{JS}(P_g(s), P_0(s))$ between the population-level step-length distribution at generation g and the initial distribution. Right: final energetic and coverage efficiencies, η_E and η_C, for the best evolved agent in each environment. Vertical markers denote the operational plateau generation estimated from a rolling-slope criterion applied to fitness and genome redistribution.

Algorithm 1 Generational evolutionary optimization of foraging strategies

Require: Population size N_{pop}, number of generations n_{gen}, genome length G, elitism fraction e_{elite}, mutation count n_{mut}, mutation scale σ_{mut}
Ensure: Final evolved population $\mathcal{P}^{(n_{\mathrm{gen}})}$

Randomly initialize the population $\mathcal{P}^{(0)} = \{X_i^{(0)}\}_{i=1}^{N_{\mathrm{pop}}}$
for $g = 0$ to $n_{\mathrm{gen}} - 1$ **do**
  for each individual $X_i^{(g)} \in \mathcal{P}^{(g)}$ **do**
    Simulate its lifetime trajectory in the current environment
    Compute fitness $\mathcal{F}_i^{(g)}$ according to Eq. (3)
  end for
  Rank all individuals by fitness
  Copy the top $N_{\mathrm{elite}} = \lfloor e_{\mathrm{elite}} N_{\mathrm{pop}} \rfloor$ individuals into $\mathcal{E}^{(g)}$
  Set parent-selection probabilities p_i according to Eq. (8)
  Initialize the offspring set $\mathcal{O}^{(g)} = \varnothing$
  while $|\mathcal{E}^{(g)}| + |\mathcal{O}^{(g)}| < N_{\mathrm{pop}}$ **do**
    Sample two parents from $\mathcal{P}^{(g)}$ using probabilities $\{p_i\}$
    Generate offspring by independent trait-wise recombination as in Eq. (9)
    Apply bounded Gaussian integer mutation as in Eq. (10)
    Add the resulting offspring to $\mathcal{O}^{(g)}$
  end while
  Form the next generation: $\mathcal{P}^{(g+1)} = \mathcal{E}^{(g)} \cup \mathcal{O}^{(g)}_{[1:N_{\mathrm{pop}} - N_{\mathrm{elite}}]}$
end for
return $\mathcal{P}^{(n_{\mathrm{gen}})}$

At each generation, all individuals are evaluated independently in the current environment and ranked by fitness. The top fraction e_{elite} is copied unchanged into the next generation, implementing elitism. In the reported experiments, $e_{\mathrm{elite}} = 0.1$, so the elite count is $N_{\mathrm{elite}} = \lfloor e_{\mathrm{elite}} N_{\mathrm{pop}} \rfloor$.

The remaining individuals are generated from parents sampled with probability proportional to fitness. If $\mathcal{F}_i$ denotes the fitness of individual i, then the selection probability is

$$p_i = \frac{\max(\mathcal{F}_i, \varepsilon_s)}{\sum_{j=1}^{N_{\mathrm{pop}}} \max(\mathcal{F}_j, \varepsilon_s)}, \tag{8}$$

where $\varepsilon_s > 0$ is a small numerical constant introduced to avoid vanishing probabilities when fitness values are extremely small.

Recombination is applied independently to each genome type. Because the genomes are interpreted as unordered lists of admissible trait values rather than ordered chromosomal sequences, each parental

genome is first randomly permuted before crossover. Let $\mathbf{a}$ and $\mathbf{b}$ denote two parental genomes of the same type, each of length G, and let π_a and π_b denote independent random permutations of $\{1, \ldots, G\}$. Writing $\mathbf{a}^{\pi_a} = (a_1^{\pi_a}, \ldots, a_G^{\pi_a})$, $\mathbf{b}^{\pi_b} = (b_1^{\pi_b}, \ldots, b_G^{\pi_b})$, the offspring genomes are formed by exchanging the second half after permutation:

$$\mathbf{c}_1 = \left(a_1^{\pi_a}, \ldots, a_m^{\pi_a}, b_{m+1}^{\pi_b}, \ldots, b_G^{\pi_b}\right), \qquad \mathbf{c}_2 = \left(b_1^{\pi_b}, \ldots, b_m^{\pi_b}, a_{m+1}^{\pi_a}, \ldots, a_G^{\pi_a}\right), \tag{9}$$

where $m = \lfloor G/2 \rfloor$. This operator mixes parental trait frequencies while avoiding an artificial dependence on locus order.

Mutation is then applied independently to each of the three genomes. In the reported experiments, exactly $n_{\mathrm{mut}} = 6$ loci were selected uniformly at random within each genome of length $G = 30$. If a_i is the value at a selected locus, the mutated value is

$$a_i' = \mathrm{clip}(a_i + \mathrm{round}(Z),\, a_{\min},\, a_{\max}), \qquad Z \sim \mathcal{N}(0, \sigma_{\mathrm{mut}}^2), \tag{10}$$

with $\sigma_{\mathrm{mut}} = 2$. The clipping operation enforces the admissible bounds of the corresponding genome: $1 \le s \le \lfloor L/2 \rfloor$ for move lengths, $1 \le v \le L$ for velocities, and $0 \le \kappa \le 35$ for turn-angle indices. The mutation and recombination hyperparameters were fixed after pilot exploration and then kept identical across all environments and grid sizes. The aim was not to optimize the evolutionary algorithm separately for each condition, but to use a single moderate setting that preserves heritability while maintaining sufficient exploratory variation. In particular, mutating 6 of 30 loci per genome provides a moderate perturbation rate, while the Gaussian mutation scale $\sigma_{\mathrm{mut}} = 2$ yields predominantly local integer changes. Likewise, shuffled half-and-half recombination was chosen because the genomes encode unordered trait-value lists rather than ordered genetic sequences. Keeping these choices fixed across all conditions reduces the risk that the reported differences between environments arise from retuning the evolutionary operators rather than from the ecological structure of the landscapes.

The full evolutionary cycle is summarized in Algorithm 1. At each generation, fitness is first evaluated for the entire population. Elite individuals are copied directly into the next generation, while the remainder of the population is filled by offspring generated through fitness-proportionate parent selection, recombination, and mutation. This procedure is repeated for the prescribed number of generations. Fitness evaluation is parallelized across agents, since trajectories are conditionally independent once the environment and the genomes are specified. In addition to the final population, the implementation stores generation-wise fitness values, summary statistics, and genome histories, enabling exact post hoc analysis of the evolutionary dynamics.

Implementation and compute resources

All simulations were executed in a Python 3.11 environment using a Docker-based GitLab CI runner. Fitness evaluation was parallelized on the CPU using a `ProcessPoolExecutor` with up to 8 worker processes. The execution host had an NVIDIA RTX A6000 GPU available (49, 140 MiB of VRAM; driver version 580.65.06, CUDA 13.0), although the simulation code itself did not use GPU-specific acceleration. A full experiment suite, comprising one complete run across all six environments for 1500 generations, required approximately 6-10 hours of wall-clock time.

B Additional robustness results

Results across different grid sizes

We tested whether the main results depend on the linear system size by repeating the analysis on 211×211 and 1009×1009 tori. These additional domains differ in both spatial extent and total resource capacity, while preserving the same resource classes, evolutionary protocol, and fitness definition used for the 503×503 system. Figures 6 and 7 show the corresponding fitness trajectories, genome-encoded step-length distributions, best evolved trajectories, and displacement-moment fits. In both additional system sizes, mean fitness increases during the early generations and then fluctuates around an environment-dependent plateau. The step-length genomes show the same qualitative evolutionary change as in the main system. Initially, broad trait lists become concentrated at short values while retaining a sparse tail of longer entries.

The evolved trajectories also preserve the same broad movement organization across system sizes. They combine local scanning with occasional relocation, although the detailed path geometry varies

Table 3: Realized resource occupancy ρ_{real}, the fraction of distinct resource-bearing sites before depletion, for the saved 503×503 maps at the five original nominal densities. Each entry describes the corresponding resource map.

Environment	$\rho = 0.005$	$\rho = 0.010$	$\rho = 0.15$	$\rho = 0.30$	$\rho = 0.45$
Uniform	0.0051	0.0100	0.150	0.299	0.449
$\mu = 1.0$	0.0049	0.0098	0.138	0.256	0.359
$\mu = 1.5$	0.0047	0.0093	0.131	0.245	0.344
$\mu = 2.0$	0.0042	0.0081	0.117	0.221	0.310
$\mu = 2.5$	0.0035	0.0068	0.097	0.184	0.263
$\mu = 3.0$	0.0028	0.0054	0.066	0.145	0.194

Table 4: Coverage-term ablation at $L = 503$ and $\rho = 0.15$. Values are mean $\pm$ standard deviation over the final 20 best-agent generations. They represent correlated late-generation best-agent evaluations from one evolutionary run per objective, rather than independent evolutionary replicates.

Environment	η_C, baseline $F = \eta_E \eta_C$	η_C, ablation $F = \eta_E$	Retained η_C (%)	Γ, ablation $F = \eta_E$
Uniform	0.9278 ± 0.0065	0.9283 ± 0.0048	100.1%	0.0190 ± 0.0013
$\mu = 1.0$	0.9282 ± 0.0057	0.9232 ± 0.0055	99.5%	0.0180 ± 0.0021
$\mu = 1.5$	0.8997 ± 0.0053	0.8920 ± 0.0103	99.1%	0.0216 ± 0.0034
$\mu = 2.0$	0.8585 ± 0.0083	0.8489 ± 0.0073	98.9%	0.0313 ± 0.0049
$\mu = 2.5$	0.8129 ± 0.0076	0.8105 ± 0.0112	99.7%	0.0373 ± 0.0100
$\mu = 3.0$	0.7834 ± 0.0118	0.7422 ± 0.0154	94.7%	0.0509 ± 0.0097

with resource environment and domain size. Larger domains allow longer realized transfers, while smaller domains impose stronger spatial constraints through the finite torus. These differences affect quantitative details but do not change the main qualitative pattern.

The bottom panels of Figures 6 and 7 compare empirical second and fourth displacement moments with the corresponding IS and LW fits. The same model-selection pattern observed on the 503×503 grid persists at both additional sizes. Across all environments and grid sizes, the score difference $\Gamma = \bar{R}^2_{\mathrm{IS}} - \bar{R}^2_{\mathrm{LW}}$ remains positive, as reported in Table 1. Thus, the preference for IS dynamics is not restricted to a single spatial scale.

Results across different resource densities

To assess the effect of resource density, we repeated the 503×503 experiments at $\rho = 0.005$, 0.010, 0.15, 0.30, and 0.45, keeping the evolutionary protocol and resource-environment classes fixed. Figure 8 summarizes the resulting best evolved trajectories, final-population step-length distributions, and the model-selection score Γ. Across all densities and environments, Γ remained positive, indicating that the intermittent-search model more accurately describes the evolved trajectories than the Lévy-walk model. The trajectory panels show that increasing density changes the spatial organization of movement, while the step-length panels show corresponding density-dependent shifts in the evolved genome. In particular, higher densities tend to concentrate step-length distributions more strongly at short values, although the strength of this effect depends on the underlying resource environment.

Fixed-genome cue-off evaluation and coverage-term ablation To assess how the sensory cue affects the expressed trajectories, we performed an evaluation-only ablation. Genomes evolved under the standard $p_{\mathrm{cue}} = 0.5$ condition were held fixed and re-evaluated with $p_{\mathrm{cue}} = 0$, without further evolution. Cue-on and cue-off evaluations used matched generations and identical resource fields. Removing the cue reduced Γ in the uniform, $\mu = 1.0$, $\mu = 1.5$, and $\mu = 2.0$ environments, but increased it for $\mu = 2.5$ and $\mu = 3.0$. Thus, the cue affects the model-comparison margin in an environment-dependent manner. Because the genomes were not re-evolved without the cue, this evaluation does not establish that IS-like dynamics evolve independently of sensing.

To test whether the coverage term drives the model classification, we re-evolved the agents using $F = \eta_E$, while retaining the remaining simulation and evolutionary settings. Removing η_C from the objective did not eliminate high coverage efficiency or the positive IS–LW fit difference. The

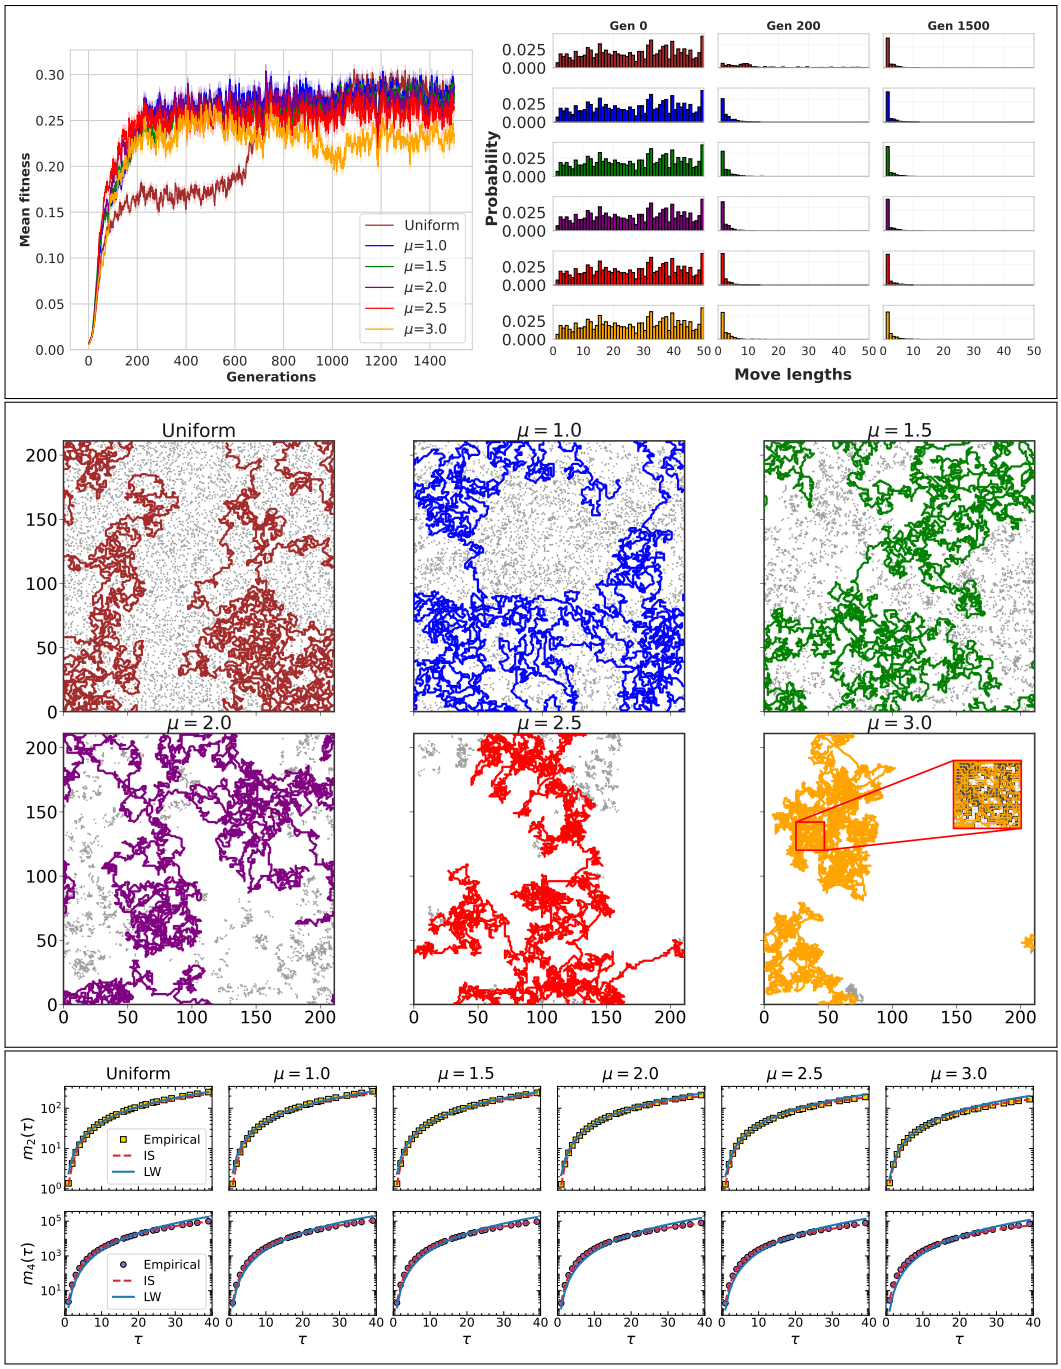


Figure 6: Additional-size results for the 211×211 torus. **Top left**: mean fitness over 1500 generations for each resource environment, with shaded standard errors across agents within the single evolutionary run at each generation. **Top right**: distributions of genome-encoded step-length entries at generations 0, 200, and 1500. **Middle**: resource maps overlaid with the trajectory of the best evolved agent in each environment. **Bottom**: empirical second and fourth displacement moments compared with best-fit IS and LW models.

coverage term may still affect the degree of redundancy, particularly for $\mu = 3.0$, but it is not solely responsible for the classification result.

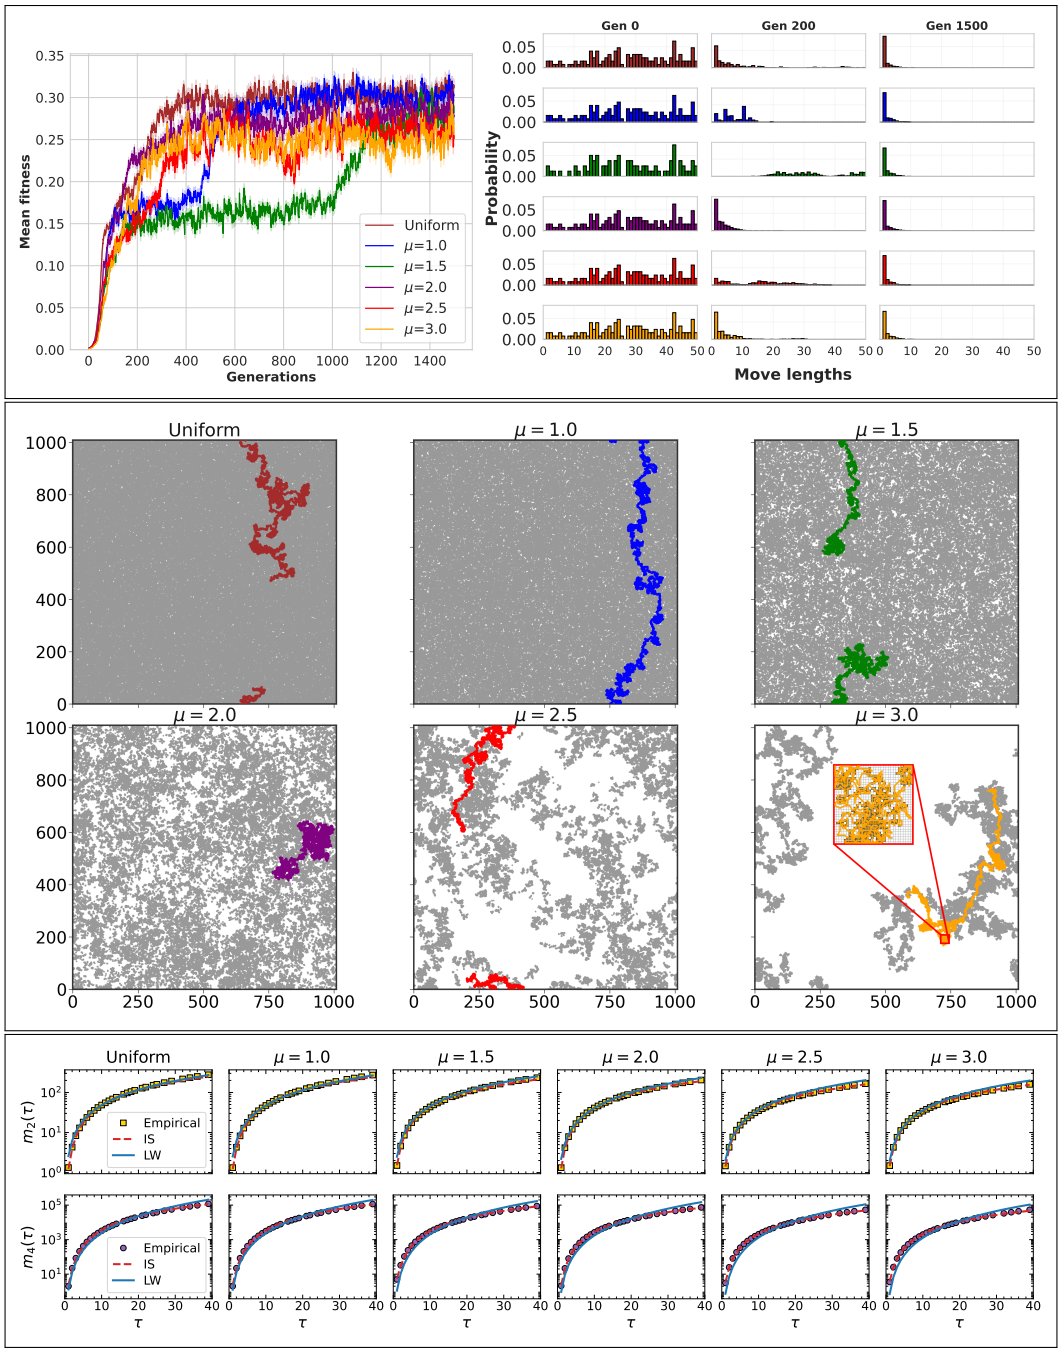


Figure 7: Additional-size results for the 1009×1009 torus. **Top left**: mean fitness over 1500 generations for each resource environment, with shaded standard errors across agents within the single evolutionary run at each generation. **Top right**: distributions of genome-encoded step-length entries at generations 0, 200, and 1500. **Middle**: resource maps overlaid with the trajectory of the best evolved agent in each environment. **Bottom**: empirical second and fourth displacement moments compared with best-fit IS and LW models.

Synthetic validation of the classifier

Figure 9 reports classification results for synthetic LW and IS trajectories with known generating models, together with parameter-recovery results for the IS model. The parameter-recovery coefficients of determination range from 0.78 to 0.95, with fitted slopes between 1.00 and 1.06.

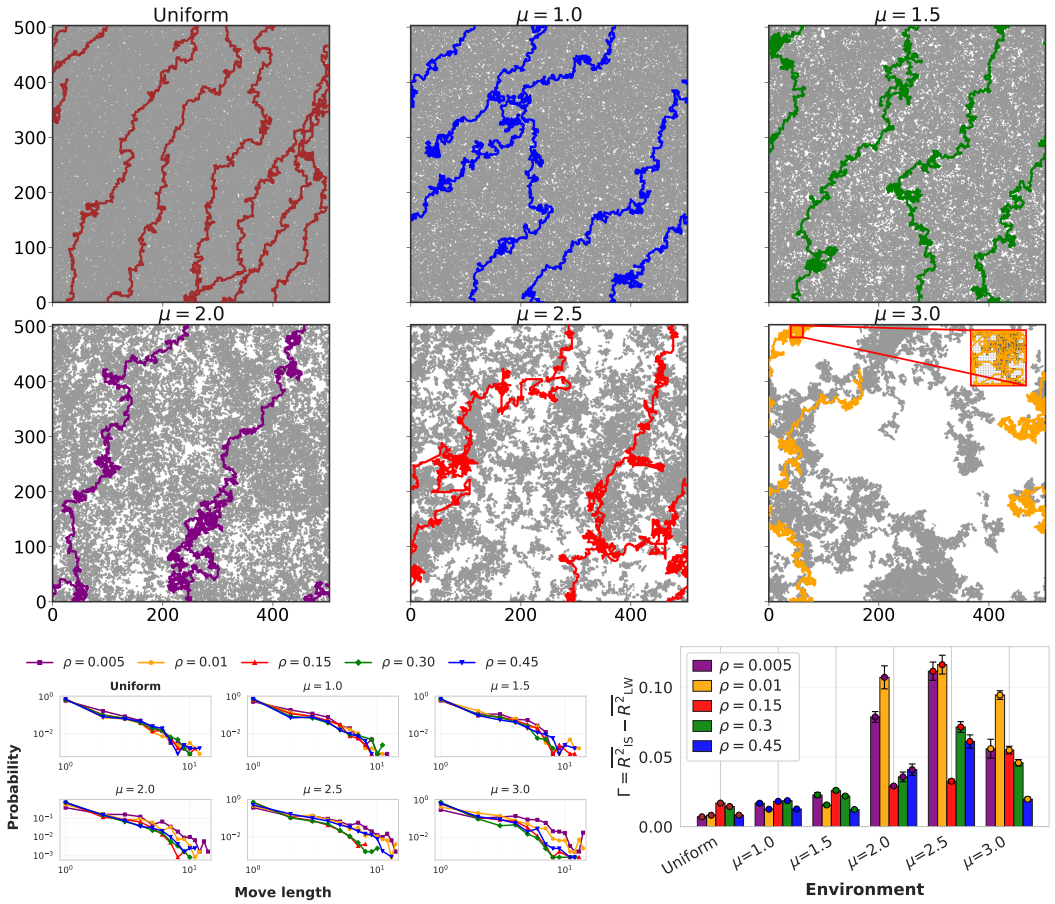


Figure 8: Robustness of evolved search dynamics across resource densities on the 503×503 torus. **Top**: best evolved trajectories overlaid on the corresponding resource maps for the highest tested density, $\rho = 0.45$. **Bottom left**: final-population distributions of genome-encoded step lengths for $\rho \in \{0.005, 0.010, 0.15, 0.30, 0.45\}$, shown separately for each resource environment. **Bottom right**: model-selection score Γ across densities, with error bars showing standard error over the final 500 generations.

Temporal structure of evolved search trajectories

To quantify temporal organization beyond trajectory visualization, we computed three waiting-time observables from the best evolved trajectories. Let $\mathbf{r}_t \in \mathbb{Z}^2$, $t = 0, \ldots, T_{\text{life}}$, denote the lattice position of an agent during one lifetime. For a visited site $\mathbf{z}$, define the ordered set of visit times $\mathcal{V}(\mathbf{z}) = \{t : \mathbf{r}_t = \mathbf{z}\} = \{t_1^{(\mathbf{z})} < \cdots < t_{m_\mathbf{z}}^{(\mathbf{z})}\}$. Return intervals are the successive gaps between visits to the same site, $\Delta t_{\mathbf{z},k}^{\text{return}} = t_{k+1}^{(\mathbf{z})} - t_k^{(\mathbf{z})}$, $k = 1, \ldots, m_\mathbf{z} - 1$, pooled over all sites with $m_\mathbf{z} \geq 2$. These intervals quantify recurrence to already visited locations.

If $h_1 < \cdots < h_{N_{\text{hit}}}$ are the resource encounter times, resource-hit intervals are defined as $\Delta t_j^{\text{hit}} = h_{j+1} - h_j$, $j = 1, \ldots, N_{\text{hit}} - 1$. These intervals quantify the temporal spacing of successful foraging events. Finally, let t_n^{new} be the first time at which the n-th distinct lattice site has been discovered, $t_n^{\text{new}} = \min\{t : |\{\mathbf{r}_0, \ldots, \mathbf{r}_t\}| = n\}$. The new site discovery interval is then $\tau_n = t_{n+1}^{\text{new}} - t_n^{\text{new}}$. This definition follows the visitation-statistics convention in which τ_n measures the waiting time to discover a new site after n distinct sites have already been visited. It is therefore distinct from return intervals, which measure recurrence to previously visited sites. For each waiting-time observable Y,

Table 5: Fixed-genome cue-off evaluation at $L = 503$ and $\rho = 0.15$. Genomes evolved under the standard $p_{\text{cue}} = 0.5$ condition were re-evaluated with $p_{\text{cue}} = 0$, without further evolution. The cue-on and cue-off Γ values are the mean $\pm$ standard deviation over the final 500 matched best-agent generations. The final column reports the mean adjusted LW coefficient of determination in the cue-off condition. These results are separate from the five-run aggregates in Table 1.

Environment	$\Gamma_{\text{cue on}}$	$\Gamma_{\text{cue off}}$	$\bar{R}^2_{\text{LW,cue off}}$
Uniform	0.0175 ± 0.0014	0.0010 ± 0.0001	0.9988
$\mu = 1.0$	0.0185 ± 0.0017	0.0029 ± 0.0008	0.9969
$\mu = 1.5$	0.0204 ± 0.0024	0.0059 ± 0.0024	0.9933
$\mu = 2.0$	0.0326 ± 0.0046	0.0145 ± 0.0062	0.9818
$\mu = 2.5$	0.0251 ± 0.0060	0.0267 ± 0.0175	0.9720
$\mu = 3.0$	0.0424 ± 0.0085	0.0797 ± 0.0483	0.9181

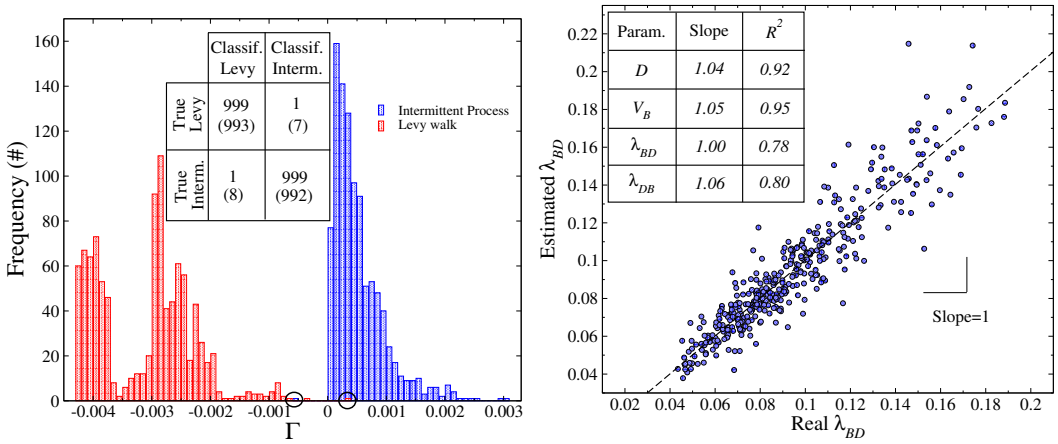


Figure 9: Synthetic validation of the moment-based classifier and intermittent-search parameter estimation. Left: distributions of the model-comparison score $\Gamma = \bar{R}^2_{\text{IS}} - \bar{R}^2_{\text{LW}}$ for 1000 synthetic Lévy-walk trajectories and 1000 synthetic intermittent-search trajectories. Positive values favor IS, whereas negative values favor LW. The confusion matrix shows that 999 of 1000 trajectories from each generating model were classified correctly. Right: estimated versus generating values of λ_{BD} for synthetic IS trajectories. The dashed line marks exact recovery. The inset summarizes the regression slopes and coefficients of determination for D, v_B, λ_{BD}, and λ_{DB}.

with observed values, $Y_1, \ldots, Y_{n_Y}$, we computed the empirical survival function

$$\widehat{S}_Y(y) = \frac{1}{n_Y} \sum_{j=1}^{n_Y} \mathbf{1}(Y_j \geq y) \, . \tag{11}$$

Thus, $\widehat{S}_Y(y)$ is the fraction of observed intervals whose duration is at least y. Slowly decaying survival curves indicate broad upper tails, corresponding to many short intervals together with rarer long gaps. We summarized this upper-tail spread by $\frac{q_{0.95}(Y)}{\text{median}(Y)}$, where $q_{0.95}(Y)$ is the empirical 95th percentile.

Table 6 summarizes these quantities after aggregation across the three grid sizes. Return intervals show the strongest right-skewed structure, resource-hit intervals are moderately right-skewed, and new-site discovery intervals remain comparatively narrow. This distinction is important: the strongest intermittency is observed in recurrence and resource-encounter timing, whereas spatial discovery proceeds at a comparatively regular rate.

Figure 10 clarifies how temporal intermittency differs across observables. Since the survival function $\widehat{S}_Y(y)$ gives the fraction of intervals with duration at least y, more slowly decaying curves indicate broader upper tails and a greater incidence of rare long intervals. Return intervals show the broadest tails, resource-hit intervals show intermediate tails, and new-site discovery intervals are comparatively narrow. This indicates that the strongest intermittency in the evolved trajectories appears in recurrence

Table 6: Waiting-time statistics for the best evolved agents. Fitness and efficiency values - cf. Eqs. (2) and (3) - are reported as mean ± standard deviation across grid sizes . Upper-tail spread is summarized by the median across grid sizes of $q_{0.95}/\mathrm{median}$. Return intervals quantify recurrence to previously visited sites, resource-hit intervals quantify the temporal spacing between resource encounters, and new-site discovery intervals quantify the timing of first visits to previously unvisited sites.

Env.	$\mathcal{F}$	η_E	η_C	Return $q_{0.95}/\mathrm{mdn.}$	Hit $q_{0.95}/\mathrm{mdn.}$	New-site $q_{0.95}/\mathrm{mdn.}$
$\mu = 1.0$	0.40 ± 0.02	0.43 ± 0.01	0.92 ± 0.02	25.4	3.0	2.0
$\mu = 1.5$	0.39 ± 0.01	0.43 ± 0.01	0.90 ± 0.01	18.6	3.5	2.0
$\mu = 2.0$	0.40 ± 0.01	0.46 ± 0.01	0.86 ± 0.01	81.5	6.0	2.0
$\mu = 2.5$	0.37 ± 0.02	0.48 ± 0.02	0.81 ± 0.01	55.6	6.0	2.1
$\mu = 3.0$	0.37 ± 0.05	0.49 ± 0.04	0.76 ± 0.06	42.9	5.0	3.0
Uniform	0.40 ± 0.02	0.43 ± 0.02	0.93 ± 0.02	39.7	3.0	1.0

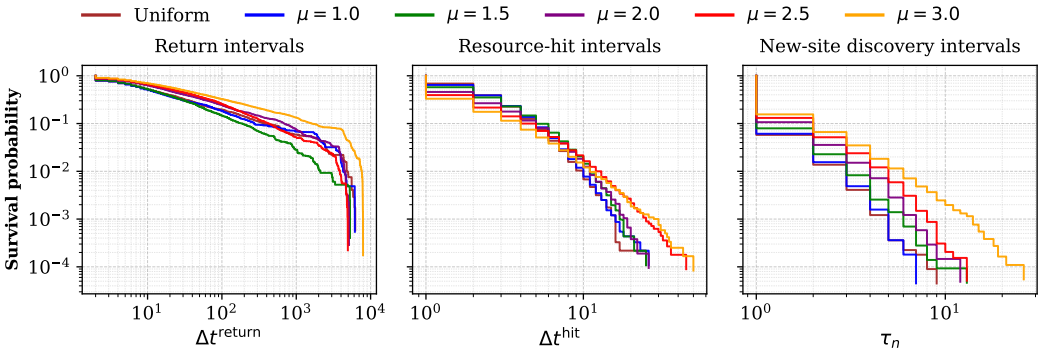


Figure 10: Empirical survival distributions of waiting-time observables in evolved search trajectories. **Left**: return intervals $\Delta t^{\mathrm{return}}$, measuring the time between successive visits to the same previously visited site. **Middle**: resource-hit intervals Δt^{hit}, measuring the time between successive resource encounters. **Right**: new-site discovery intervals τ_n, measuring the waiting time between first visits to successive previously unvisited sites.

to previously visited sites and in the timing of resource encounters, whereas discovery of new sites proceeds on a more regular timescale.

C Analytical expressions for the second and fourth moments of an intermittent process

To select between Lévy and Intermittent strategies, we use the derivations published recently in [32] and implemented as an open source Python library (cf. [33]). These results enable us to distinguish between a planar uniform LW process, as described in Section 2.4 (cf. Eq. (5)), and an IS model.

As mentioned in the main text, for the IS model, we can derive the analytical expressions of the second and the fourth moments. Namely, the second moment is given by [32]:

$$\langle \mathbf{r}^2(t) \rangle = C_1^{(2)} t - C_2^{(2)} (1 - e^{-\beta\alpha t}) \tag{12}$$

with the values of the constants given by

$$\begin{aligned} C_1^{(2)} &= 2\frac{1-\alpha}{\alpha} D_I + 4\alpha D, \\ C_2^{(2)} &= 2\frac{1-\alpha}{\alpha^2 \beta} D_I, \end{aligned}$$

where $\beta = \lambda_{BD} + \lambda_{DB}$, $\alpha = \lambda_{BD}/\beta$, and $D_I = v_B^2/\beta$, with λ_{BD} and λ_{DB} the transition rates from the ballistic to the diffusive regime and vice-versa, respectively.

The fourth moment is given by

$$\langle \mathbf{r}^4(t) \rangle = C_1^{(4)} t^2 + C_2^{(4)} t - C_3^{(4)} - C_4^{(4)} e^{-\beta t} + \left(C_5^{(4)} t^2 + C_6^{(4)} t - C_7^{(4)} \right) e^{-\beta \alpha t} \quad (13)$$

with

$$\begin{aligned}
C_1^{(4)} &= 8\alpha^2 \left(2D + \frac{1-\alpha}{\alpha^2} D_I \right)^2, \\
C_2^{(4)} &= 8 \frac{1-\alpha}{\alpha^3 \beta^2} \left[\left(D - \frac{\alpha+1}{2\alpha^2} D_I \right)^2 - \frac{5}{8\alpha^4} D_I^2 \right], \\
C_3^{(4)} &= 4 \frac{1-\alpha}{\alpha^4 \beta^4} \left[\left(2D - \frac{\alpha^2+\alpha+1}{\alpha^3} D_I \right)^2 - \frac{6\alpha^3 - 6\alpha^2 - \alpha - 2}{2\alpha^6} \right], \\
C_4^{(4)} &= 32(1-\alpha)\alpha\beta \left(2D + \frac{1}{1-\alpha} D_I \right)^2, \\
C_5^{(4)} &= 4 \frac{1-\alpha}{\alpha^2} D_I^2, \\
C_6^{(4)} &= \frac{16}{\alpha^2 \beta} D_I^2, \\
C_7^{(4)} &= \frac{8}{\alpha^2 \beta^2} D_I \left(8D + \frac{11\alpha^2 - 12\alpha + 3}{\alpha^2(1-\alpha)} D_I \right)^2.
\end{aligned}$$

These are the expressions of the moments used to compute m_k^{mod} in Eq. (6) for the IS model.

The derivation of these analytical expressions is given in detail in Ref. [32] and can be summarized as follows:

- Main idea: The IS process is written as a two-state planar stochastic process, where a walker alternates between the ballistic phase (phase "+" below), with speed v_B, and a diffusive phase D (phase "−" below), with diffusivity D.
- Main aim: derive an analytical expression for the probability density function for each position $\mathbf{r}$ and time t, expressed as the sum of densities defined for each phase, namely

$$P(\mathbf{r}, t) = P_+(\mathbf{r}, t) + P_-(\mathbf{r}, t)$$

  .
- One then introduce auxiliary quantities, namely densities $\nu_\pm(\mathbf{r}, t)$ describing the switching between phases, residence-time densities in each phase, given by $\psi_\pm(t)$, and conditional displacement densities during each phase, denoted by $w_\pm(\mathbf{r}, t)$.
- The switching densities $\nu_\pm(\mathbf{r}, t)$ describe the flux of trajectories that arrive at $(\mathbf{r}, t)$ while switching from the opposite phase into phase $\pm$. With initial phase probabilities $c_\pm$, the renewal equations take the form $\nu_\pm(\mathbf{r}, t) = c_\mp q_\mp(\mathbf{r}, t) + \int_0^t \mathrm{d}t' \int_{\mathbb{R}^2} \mathrm{d}\mathbf{r}'\, q_\mp(\mathbf{r}', t') \nu_\mp(\mathbf{r} - \mathbf{r}', t - t')$, and $P_\pm(\mathbf{r}, t) = c_\pm Q_\pm(\mathbf{r}, t) + \int_0^t \mathrm{d}t' \int_{\mathbb{R}^2} \mathrm{d}\mathbf{r}'\, Q_\pm(\mathbf{r}', t') \nu_\pm(\mathbf{r} - \mathbf{r}', t - t')$, where $q_\pm(\mathbf{r}, t) = w_\pm(\mathbf{r}, t)\psi_\pm(t)$, and $Q_\pm(\mathbf{r}, t) = w_\pm(\mathbf{r}, t) \int_t^\infty \psi_\pm(t')\, \mathrm{d}t'$.
- The renewal equations are then transformed into the Fourier–Laplace space, where the convolutions become products. After algebraic manipulation, the transformed probability density function for each position and time reads

$$P(\mathbf{k}, s) = \frac{Q_+(\mathbf{k}, s)\,[c_+ + c_- q_-(\mathbf{k}, s)] + Q_-(\mathbf{k}, s)\,[c_- + c_+ q_+(\mathbf{k}, s)]}{1 - q_+(\mathbf{k}, s) q_-(\mathbf{k}, s)}.$$

- The model is then specified by exponential residence-time densities, $\psi_+(t) = \lambda_{BD} e^{-\lambda_{BD} t}$, $\psi_-(t) = \lambda_{DB} e^{-\lambda_{DB} t}$, where λ_{BD} and λ_{DB} are the transition rates from

phase $+$ to phase $-$ and from $-$ to $+$ respectively. Substituting these expressions in the functions $q_{\pm}(\mathbf{k}, s)$ and $Q_{\pm}(\mathbf{k}, s)$, one obtains after algebraic manipulation

$$P(\mathbf{k}, s) = \frac{c_B(s + \lambda_{DB} + Dk^2) + c_D\sqrt{(s + \lambda_{BD})^2 + v_B^2 k^2} + c_B\lambda_{BD} + c_D\lambda_{DB}}{(s + \lambda_{DB} + Dk^2)\sqrt{(s + \lambda_{BD})^2 + v_B^2 k^2} - \lambda_{BD}\lambda_{DB}},$$

with $c_B = \frac{\lambda_{DB}}{\lambda_{BD}+\lambda_{DB}}$ and $c_D = \frac{\lambda_{BD}}{\lambda_{BD}+\lambda_{DB}}$.

- Finally, the second and fourth moments are obtained by differentiating this propagator at $\mathbf{k} = \mathbf{0}$ and then applying the inverse Laplace transform $\mathcal{L}^{-1}$:

$$\langle \mathbf{r}^2(t) \rangle = \mathcal{L}^{-1}\left[-\left(\frac{\partial^2}{\partial k_x^2} + \frac{\partial^2}{\partial k_y^2} \right) P(\mathbf{k}, s)\bigg|_{\mathbf{k}=\mathbf{0}} \right],$$

and

$$\langle \mathbf{r}^4(t) \rangle = \mathcal{L}^{-1}\left[\left(\frac{\partial^4}{\partial k_x^4} + 2\frac{\partial^4}{\partial k_x^2 \partial k_y^2} + \frac{\partial^4}{\partial k_y^4} \right) P(\mathbf{k}, s)\bigg|_{\mathbf{k}=\mathbf{0}} \right].$$

For the LW model, since the velocity v_L enters the definitions of the moments only trivially, it is used here as a scaling parameter. The two remaining parameters, τ_0 and γ, are estimated with a grid search algorithm. Further details on the derivations and their implementation in the *IntLevPy* library are given in Refs. [32, 33].

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction state that the paper studies the evolution of foraging strategies in finite depleting landscapes and shows that the evolved trajectories are better described by intermittent-search dynamics than by strict Lévy walks. These claims match the simulation results and the moment-based model comparison reported in Secs. 2–4 and Table 1.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: The discussion identifies the main limitations of the study, including finite periodic domains, finite lifetimes, non-renewable resources, fixed evolutionary hyperparameters, and single-agent evaluations. It also describes future directions involving larger domains, resource renewal, alternative sensing rules, and multi-agent extensions. The manuscript further states that the work is a controlled simulation study rather than a deployed autonomous-search system.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation, while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best

judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [N/A]

   Justification: The paper does not introduce new formal theorems or original analytical proofs. The analytical moment expressions used for intermittent-search fitting are taken from prior work and are cited in the manuscript. The appendix states these expressions and summarizes the previously published renewal-equation construction to clarify how the model functions used in the fitting procedure are obtained.

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: The manuscript specifies the simulation environment, resource-generation procedure, genome representation, movement rule, fitness definition, evolutionary operators, grid sizes, resource environments, and the moment-based classification pipeline. Appendix A gives the generational evolutionary algorithm, parameter table, selection rule, recombination rule, mutation operator, and compute setup. The GitHub repository provides the simulation and post hoc analysis code, together with the generated result archives for the reported grid sizes through Git LFS. These archives contain the trajectories, genome histories, fitness records, resource maps, and intermediate outputs needed to reproduce the manuscript figures and summary analyses without rerunning all simulations. The same outputs can also be regenerated from the released code using the reported parameter settings.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.

- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: The source code and generated simulation outputs for the evolutionary foraging simulations and post hoc analyses are available through the public repository:`https://github.com/shailendrabhandari/intermittent-vs-levy-foraging.git`. The repository contains the simulation code, analysis workflow, figure-generation notebook, documentation, and Git LFS archives with the generated outputs used for the reported analyses. These archives include the trajectories, genome histories, fitness records, resource maps, final populations, and intermediate post hoc outputs for the reported grid sizes. The main results can therefore be reproduced either by running the analysis workflow on the provided archives or by regenerating the outputs from the released simulation code using the parameter settings reported in the manuscript.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes]

Justification: The manuscript provides the experimental details relevant to the evolutionary simulations, including resource-landscape construction, genome encoding, population size, number of generations, lifetime, grid sizes, density settings, elitism fraction, mutation count, mutation scale, movement cost, resource gain, boundary condition, and random seed. The model-comparison procedure, including empirical moment computation, fitted IS and LW model classes, log-moment loss, and Γ score, is described in the Methods. Additional implementation details are provided in the Appendix.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes].

Justification: The reported experimental results include uncertainty estimates for the quantities supporting the main claims. Fitness trajectories are plotted with shaded standard errors, and the model-comparison table reports mean values with associated standard deviations for Γ, $\overline{R}^2_{\mathrm{IS}}$, and $\overline{R}^2_{\mathrm{LW}}$. The relevant figure and table captions specify whether the uncertainty shown is a standard error or a standard deviation, and the model-comparison table states the analysis unit over which the reported variation is computed.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Appendix includes an "Implementation and compute resources" paragraph that reports the Python environment, Docker-based runner, CPU parallelization strategy, available hardware, and approximate wall-clock runtime for a full experiment suite. It also clarifies that the simulation code uses CPU parallelism rather than GPU-specific acceleration.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The work consists of computational simulations of foraging dynamics and does not involve deceptive practices, unsafe data collection, or human-subject experimentation. We are not aware of any aspect of the research that conflicts with the NeurIPS Code of Ethics.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: The Discussion briefly identifies potential positive uses in energy-aware autonomous exploration and ecological monitoring, while also noting that applications in sensitive human or environmental settings would require privacy, safety, and oversight safeguards. The manuscript further clarifies that the present contribution is a controlled simulation study rather than a deployed system.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.

- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: The work does not release or rely on high-risk generative models, scraped datasets, or sensitive data assets. The contribution is a simulation framework for evolutionary search in abstract resource landscapes.

    Guidelines:

    - The answer [N/A] means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes].

    Justification: The paper cites the external literature and software components used for moment-based intermittent-search versus Lévy-walk classification, including the IntLevPy analysis library. No third-party empirical datasets, pretrained models, or proprietary assets are used in the simulations. The released repository includes license information for the new code and cites the external methodological sources on which the classification workflow is based.

    Guidelines:

    - The answer [N/A] means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
    - If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented, and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: The paper releases new simulation and analysis code for the evolutionary foraging experiments. The GitHub repository documents the simulation script, analysis workflow, requirements, model parameters, and reproduction procedure. The generated outputs can be regenerated from the released code and the parameter settings reported in the manuscript.

Guidelines:

- The answer [N/A] means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [N/A]

Justification: The paper does not involve crowdsourcing or research with human subjects. All results are obtained from computational simulations.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [N/A]

Justification: The paper does not involve human subjects, participant recruitment, or any human-data collection requiring IRB review.

Guidelines:

- The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: Large language models are not part of the core methodology, experiments, or scientific claims of this work.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.