

Collective search with finite perception: transient dynamics and search efficiency

Adam Gosztolai*, Jose A. Carrillo, and Mauricio Barahona†
Department of Mathematics, Imperial College London, United Kingdom
(Dated: November 29, 2018)

Motile organisms often use finite spatial perception of their surroundings to navigate and search their habitats. Yet standard models of search are usually based on purely local sensory information. To model how a finite perceptual horizon affects ecological search, we propose a framework for optimal navigation that combines concepts from random walks and optimal control theory. We show that, while local strategies are optimal on asymptotically long and short search times, finite perception yields faster convergence and increased search efficiency over transient time scales relevant in biological systems. The benefit of the finite horizon can be maintained by the searchers tuning their response sensitivity to the length scale of the stimulant in the environment, and is enhanced when the agents interact as a result of increased consensus within subpopulations. Our framework sheds light on the role of spatial perception and transients in search movement and collective sensing of the environment.

INTRODUCTION

Exploration, movement, and search for resources are ubiquitous among organisms in nature [1–3]. Classical theories of search [4], such as optimal foraging theory [5, 6], have mostly focused on long time limits and typically assume that natural selection favours search strategies that maximise long-term encounters with nutrients. However, many phenomena in ecology [7] and other fields of biology operate in transient regimes [8, 9], extending over time scales that never reach the asymptotic stationary state [10]. Another typical assumption is to consider random walks [11–13] or diffusion processes [14] to describe the movement of searchers navigating the landscape based on local information [15–17]. Yet, in many instances, searchers can obtain and store [18] non-local information gathered through sensory cues [19, 20] or through anticipation of environmental changes [21–23] (Fig. 1). The question then arises as to how such finite perceptual range can influence both the dynamics of movement and the search efficiency over the finite time scales relevant in biology.

Here we study the role of finite time scales associated with ecological movement and search; specifically, the effect of limited spatial perception when the search time is itself finite. To formalise these aspects, we propose an *optimal navigation* (ON) model, which allows us to extend the description of search as a biased random walk [11, 14, 24], and reinterpret it in the framework of optimal control theory. The ON model includes a time horizon that quantifies the perceptual range of the searchers along their trajectory and fixes a non-local optimisation target for the agents. In the limit of vanishing time horizon (i.e., as the spatial perception shrinks and the information becomes local), the ON model recovers the classic Keller-Segel (KS) drift-diffusion model [25]

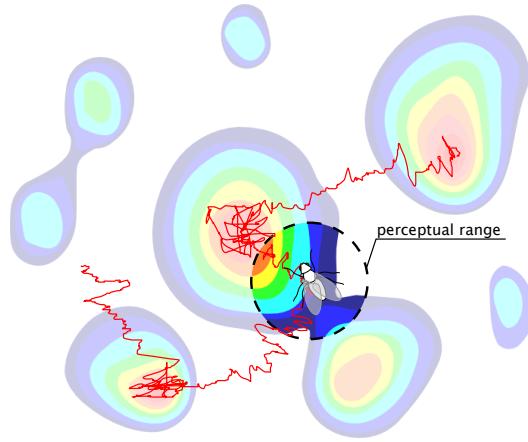


Figure 1. A searcher with a finite perceptual range navigating a heterogeneous landscape using a biased random walk search strategy. In contrast to standard local searchers, which navigate based only on point-wise information, our searcher can use non-local information within its perceptual range to optimise its movement and exploration.

of *local* search strategies (i.e., with instantaneous sensing and alignment to the point-wise gradient).

Using simulations and analytical results, we find that a population of non-local searchers moving towards a nutrient patch exhibits distinct transient behaviour, clustering faster at the hotspot than local searchers, thereby increasing their search efficiency. Our results show that the maximum efficiency gain occurs when the perceptual range of the searchers matches the environmental length scale over which the nutrient concentration changes significantly. As the search time becomes asymptotically large or small, the efficiency gain from the non-local strategy diminishes, and the searchers behave effectively as local responders. If the environmental length scale changes, we show that the efficiency gain can be maintained as long as searchers can adjust their sensitivity. This means that finite perception remains advantageous to searchers that can rescale their response dynamically, or to populations that contain a diversity of responses. Finally,

* adam.gosztolai@unige.ch

† m.barahona@imperial.ac.uk

we consider the effect of interaction between searchers, and show that non-local information consistently reinforces the dominant strategy in the population and leads to improved search efficiency overall, even though multimodality (subpopulations) can appear during the transients. Our framework provides and optimisation perspective on a range of collective phenomena in population biology and, more generally, on biologically-inspired search and exploration algorithms, thus shedding light on the role of spatial perception on finite-time search.

I. THE CLASSIC KELLER-SEGEL MODEL: NOISY SEARCH WITH LOCAL GRADIENT ALIGNMENT

A classic model for the dynamics of a population of searchers using local gradient alignment is given by the Keller-Segel equation, which we briefly recap here.

Consider a population of searchers moving in a closed, bounded, d -dimensional domain $\Omega \subset \mathbb{R}^d$. The searchers move by responding to two concentration fields: to a primary stimulant $S_1(x)$ (e.g., nutrient) with sensitivity χ_1 , and to a secondary stimulant $S_2(x, t)$ (e.g., pheromone), released by the agents themselves, with sensitivity χ_2 .

On long space and time scales relative to the microscopic motion, one can describe the biased random walk $x(t) \in \Omega$ of searchers by a Langevin equation [24]:

$$\frac{dx}{dt} = v(x, t) + \sqrt{2D} \xi(t), \quad (1)$$

where $\xi(t)$ is a white noise process, D is the coefficient of diffusion, and

$$v(x, t) = \nabla (\chi_1 S_1(x) + \chi_2 S_2(x, t)). \quad (2)$$

is the *velocity of the searcher*. The parameters D , χ_1 and χ_2 are typically inferred experimentally from trajectories of the agents [14, 26], and can sometimes be expressed in terms of microscopic parameters [27]. Note that Eqs. (1)–(2) describe a searcher that uses *local information*, since it aligns its velocity instantaneously to the gradients of S_1 and S_2 .

The secondary stimulant S_2 introduces interaction between the searchers. If S_2 is assumed to diffuse faster than the searchers, its evolution is given by

$$S_2(x, t) = \int_{\Omega} \Phi(x - x') \rho(x', t) dx' := \Phi * \rho \quad (3)$$

where

$$\Phi(x) = \begin{cases} -\frac{1}{2\pi} \log \|x\|, & \text{for } d = 2 \\ \frac{\Gamma(\frac{d}{2} + 1)}{d(d-2)\pi^{d/2}} \|x\|^{2-d} & \text{for } d \neq 2, \end{cases} \quad (4)$$

and Γ is the Gamma function (see Appendix A).

Taking all together, the time evolution of the population density $\rho(x, t)$ of searchers obeying (1)–(3) can

be described with a Fokker-Planck equation [28] known as the *Keller-Segel* (KS) model. In dimensionless variables $x \rightarrow x/L$ and $t \rightarrow t/T_D$, where $T_D = L^2/D$ is the diffusion time of the searchers, the KS equation reads

$$\partial_t \rho - \nabla^2 \rho + \text{Pe}_1 \nabla \cdot (\rho \nabla S_1) + \text{Pe}_2 \nabla \cdot (\rho \nabla (\Phi * \rho)) = 0, \quad (5)$$

where the parameters $\text{Pe}_1 = \chi_1 S_{1,\text{av}}/D$ and $\text{Pe}_2 = \chi_2 S_{2,\text{av}}/D$ are Péclet numbers quantifying the ratio of diffusive to advective forces on the searchers, and $S_{1,\text{av}}$, $S_{2,\text{av}}$ are the average stimulant concentrations. Given an initial distribution $\rho_0(x) := \rho(x, 0)$, a stimulant profile $S_1(x)$, and parameters Pe_1 and Pe_2 , Eq. (5) can be solved numerically using standard techniques. Here, we use a first-order in time, second-order in space forward Euler scheme [29] with upwind discretisation and $\Delta x = 0.01$ and $\Delta t = 10^{-6}$. We denote the solution of the KS model by $\rho_{\text{KS}}(x, t)$. See Appendix A for a fuller derivation and details of the non-dimensionalisation.

Variational rewriting of the KS model

The KS model can be recast in a variational gradient formulation [30, 31]. First, rearrange (5) as an advection equation:

$$\partial_t \rho + \nabla \cdot (\rho u) = 0 \quad (6)$$

$$\text{with } u = -\nabla \log \rho + v, \quad (7)$$

where u is the *velocity of the population*. Since the velocity of the searchers v is a gradient (2), u is also a gradient:

$$u = \nabla (-\log \rho + \text{Pe}_1 S_1 + \text{Pe}_2 (\Phi * \rho)). \quad (8)$$

All terms in (8) are either local or symmetric with respect to x ; hence u can be written [30] as a first variation

$$u = -\nabla \left(\frac{\delta \mathcal{F}}{\delta \rho} \right), \quad (9)$$

where the *free energy* functional

$$\mathcal{F}[\rho] = \int_{\Omega} \left(\rho \log \rho - \text{Pe}_1 \rho S_1 - \frac{\text{Pe}_2}{2} \rho (\Phi * \rho) \right) dx \quad (10)$$

includes (in order): an entropic term from the stochastic component of the dynamics (2); the internal energy of the stimulant landscape S_1 ; and a term from the interaction between searchers via the secondary stimulant S_2 .

The KS model can then be rewritten in the equivalent gradient flow form

$$\partial_t \rho = \nabla \cdot \left(\rho \nabla \frac{\delta \mathcal{F}}{\delta \rho} \right), \quad (11)$$

which has the important implication that the evolution of $\rho(x, t)$ can be computed using the Jordan-Kinderlehrer-Otto (JKO) variational optimisation scheme [31, 32].

From an initial density $\rho_0(x)$, the JKO scheme constructs a sequence of probability distributions $\{\rho(x, k\Delta t)\}_{k \geq 0}$

$$\rho(x, (k+1)\Delta t) = \operatorname{argmin}_{\rho'} \left\{ \frac{1}{2\Delta t} d_W(\rho(x, k\Delta t), \rho') + \mathcal{F}[\rho'] \right\}, \quad (12)$$

where $\Delta t > 0$ is the time step, and $d_W(\cdot, \cdot)$ is the Wasserstein distance between two distributions. The solution (12) has been proved to converge to the solution of Eq. (5) in the limit $\Delta t \rightarrow 0$ [32].

II. THE OPTIMAL NAVIGATION MODEL: A POPULATION OF SEARCHERS WITH NON-LOCAL OPTIMISATION

The variational rewriting (11) and its approximation scheme (12) leads us to formulate the *optimal navigation* (ON) search model, as follows. Consider a population of searchers that move by performing the optimisation (12) over a finite time horizon $\tau \geq \Delta t > 0$, which reflects the perceptual range of the agents. Then the time evolution of the population corresponds to a sequence of constrained optimisation problems [33], i.e., a succession of JKO solutions, each over time τ .

Starting from the initial density $\rho_0(x)$, we construct the evolution of $\rho(x, t)$, such that each iteration $k \geq 0$ finds $m(x, s) := \rho(x, k\tau + s)$ for $s \in [0, \tau]$ by solving the minimisation problem:

$$\begin{aligned} \text{minimise}_{(m,u)} \mathcal{J}(m, u) &:= \int_0^\tau \int_\Omega m \frac{|u|^2}{2} dx ds + \mathcal{F}[m(x, \tau)] \\ \text{subject to } \partial_s m + \nabla \cdot (mu) &= 0. \end{aligned} \quad (13)$$

Note that the constraint is the continuity equation ensuring the conservation of ρ as in (6), whereas the cost function \mathcal{J} contains a transportation cost, which constrains the average motion to geodesics between optimal states, and an end-point term involving the evaluation of the free energy at τ (10). Although we use it here for a particular form of the free energy functional, the formulation is generic: through suitable choice of \mathcal{F} , the ON model (13) converges to a broad class of conservation laws as long as they can be recast as continuity equations and possess a variational structure [30–32, 34].

We denote the solution of the ON model (13) by $\rho_{\text{ON}}(x, t; \tau)$, and compute it using Algorithm 1, a gradient descent algorithm inspired by Ref. [35] and presented in detail in Appendix B.

Physically, the ON model (13) describes the motion of searchers that optimise their displacement over paths bounded by the time horizon τ . From the proof of the JKO scheme [32], it follows directly that the ON model recovers the local KS model as $\tau \rightarrow 0$:

$$\lim_{\tau \rightarrow 0} \rho_{\text{ON}}(x, t; \tau) = \rho_{\text{KS}}(x, t) \quad (14)$$

For finite horizon $\tau > 0$, the time evolution of the ON model departs from the KS solution due to the effect of

non-local information on the movement of the searchers, as explored below.

III. NON-LOCAL SEARCH: TRANSIENTS AND ENHANCED SEARCH EFFICIENCY

We use the ON model (13) to study how the finite perception of the agents (encapsulated in the time horizon $\tau > 0$) affects the search at the population level. We first consider non-interacting searchers insensitive to the secreted stimulant, i.e., $\text{Pe}_2 = 0$. The case of interacting searchers is presented in Section V.

Our numerics start with a uniform initial condition $\rho_0(x) = 1$ and we compute $\rho_{\text{ON}}(x, t; \tau)$, the time evolution (13) of the ON population of non-local searchers with time horizon $\tau > 0$. We also compare it to the time evolution (5) of a KS population of local searchers, or equivalently the ON model with $\tau = 0$.

Both the KS and ON models converge to the same stationary solution $\rho_\infty(x)$ asymptotically as $t \rightarrow \infty$:

$$\rho_\infty(x) := \lim_{t \rightarrow \infty} \rho_{\text{KS}}(x, t) = \frac{1}{Z} \exp(\text{Pe}_1 S_1(x)) \quad (15)$$

$$= \lim_{t \rightarrow \infty} \rho_{\text{ON}}(x, t; \tau), \quad (16)$$

for all τ and S_1 , given by the Gibbs-Boltzmann distribution (15), where Z is a normalisation constant. This result is well known for the KS equation [28]. To check that $\rho_\infty(x)$ is also the stationary solution of the ON model, note that at stationarity $d_W(\rho_\infty, \rho') = 0$ in (12), and the result (16) follows from solving for the minimiser.

The approach to stationarity, on the other hand, reveals differences between the KS and ON models. To develop our intuition, let us first consider a linear attractant gradient $S_1(x) = \alpha x$. Such a landscape is fully characterised by the local gradient $\nabla S_1(x)$, and hence non-local searchers have no advantage since $\tau > 0$ provides no further information than what is known by local searchers. In line with this expectation, when we solve the ON model (13) using the method of characteristics (Section S2 in the SI), we find that the drift velocity predicted by the KS and ON models coincide for all τ , i.e.

$$u(x, t; \tau) \equiv u(x, t) \quad \text{when } S_1 \text{ is linear.}$$

This observation implies that non-local search is only advantageous in landscapes with non-zero curvature. To illustrate this point in more detail, let us consider a static Gaussian concentration of stimulant with characteristic length scale $\sigma \ll 1$ over the domain $\Omega = [-1/2, 1/2]^d$ (Fig. 2a):

$$S_1(x; \sigma) = \frac{1}{\sqrt{(2\pi)^d \det(\Sigma)}} \exp\left(-\frac{1}{2} x^T \Sigma^{-1} x\right), \quad (17)$$

where $\det(\cdot)$ is the matrix determinant and x^T is the transpose of the vector x . For simplicity of the computations that follow, we take $\Sigma = \sigma^2 I_d$, where I_d is the d -dimensional identity matrix. Such a Gaussian landscape

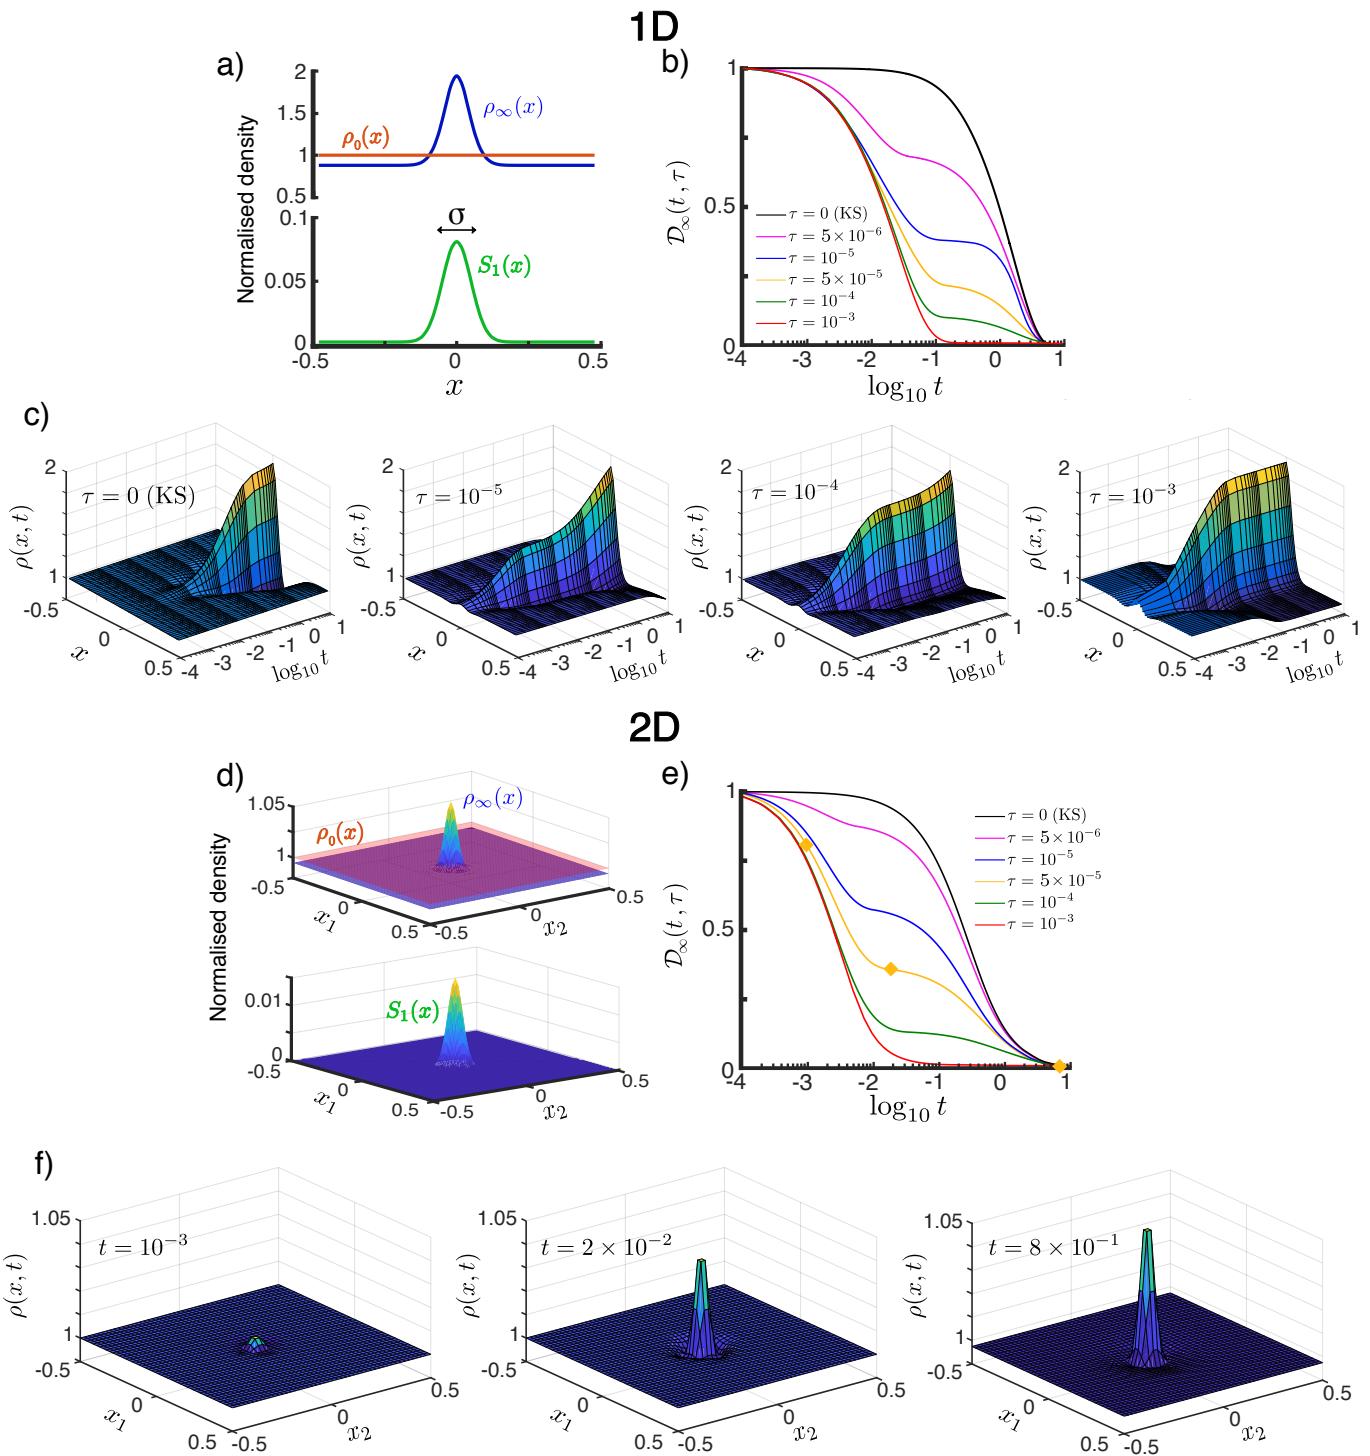


Figure 2. Transient population dynamics in nonlocal search with the ON model in the one-dimensional (a–c) and two-dimensional (d–f) cases. a) The searchers are initially uniformly distributed with $\rho_0(x) = 1$. We simulate the time evolution of the population in a Gaussian stimulant profile $S_1(x; \sigma)$ (17) until they reach the stationary state $\rho_\infty(x)$ (15). b) The approach towards stationarity measured by $\mathcal{D}_\infty(t; \tau)$ (18), the normalised L^2 -distance of the solution to the stationary state. For small time horizons $0 < \tau < \tau_\sigma^*$, there exists a long-lived intermediate transient state, whereas for $\tau \geq \tau_\sigma^*$ the dynamics directly converges to the stationary state as the searchers quickly escape the diffusion-dominated part of the domain. See Fig. 3 for the definition of τ_σ^* . c) Space-time plots illustrating the convergence towards the stationary state by the ON model in the one-dimensional case ($d = 1$). The time evolution for increasing time horizon $\tau = 0, 10^{-5}, 10^{-4}, 10^{-3}$ show qualitatively different transients. The KS model is equivalent to $\tau = 0$. d) – e) are equivalent to a) – b) but for the two-dimensional case. f) Snapshots of the time evolution of the ON population with time horizon $\tau = 5 \times 10^{-5}$ in the two-dimensional case taken at increasing times ($t = 10^{-3}, 2 \times 10^{-2}, 8 \times 10^{-1}$, marked with yellow diamonds in e)). The convergence to stationarity displays the two-stage transient: the searchers near the centre aggregate during the initial fast transient, whereas the searchers far from the centre slowly diffuse towards the centre until the stationary state is reached. As the time horizon τ is increased this second slow transient dynamics is reduced. See animations of this solution in the supporting material. All simulations in this figure with $Pe_1 = 2$ and $\sigma = 0.05$.

serves as a simple model of a stimulant patch emanating from a point source, and its characteristic length scale σ indicates regions of steep attractant gradients near the source ($\|x\| \ll \sigma$) and regions of shallow gradients in the tails of the distribution ($\|x\| \gg \sigma$).

In Fig. 2 we show that, as the time horizon τ is increased, the population of ON searchers exhibits a faster approach to stationarity, as measured by the normalised L^2 -distance between $\rho_{\text{ON}}(x, t; \tau)$ and $\rho_\infty(x)$ as a function of time for different values of τ :

$$\mathcal{D}_\infty(t; \tau) = \frac{\|\rho_{\text{ON}}(x, t; \tau) - \rho_\infty(x)\|}{\|\rho_0(x) - \rho_\infty(x)\|}. \quad (18)$$

Fig. 2a-c presents the one-dimensional case ($d = 1$). The effect of the horizon in accelerating the convergence to stationarity is shown in Fig. 2b. Note also that for small values of τ , an intermediate, quasi-stationary distribution develops during the transient (e.g., $\tau = 10^{-5}$ in Fig. 2b-c). This long-lived intermediate behaviour is the result of the population evolving on two timescales [14]: searchers near the maximum of $S_1(x)$ ($\|x\| \ll \sigma$) are driven by advection due to the steep gradient, whereas those far from the maximum ($\|x\| \gg \sigma$) are driven by diffusion in shallow gradients, and hence move more slowly towards the maximum. Due to the slow diffusive searchers, the stationary state is only reached at a longer time scale $t \sim 1$. As the horizon τ is increased, this dual behaviour (diffusion- or advection-dominated) is lost: the searchers escape quickly the diffusion-dominated part of the domain and, as a result, the distribution approaches stationarity increasingly faster with no appreciable quasi-stationary transient distribution (e.g., $\tau = 10^{-3}$ in Fig. 2b-c). The same behaviour is observed also in the two-dimensional case ($d = 2$) in Fig. 2d-f.

Such transient states can be important in biological systems, which typically operate on time scales far from the asymptotic long-time regime [7–9]. In our setting, this situation arises when the *search time* T (which is analogous to the foraging effort in ecology) is smaller than the diffusion-dominated convergence time, i.e., when $T \ll 1$. In such a situation, non-local (ON) searchers have an advantage over local (KS) searchers since they converge faster to areas with high concentration of attractant. To quantify this effect, we consider the amount of stimulant S_1 encountered over the search time T

$$\hat{U}(\tau, T) = \int_0^T \int_\Omega S_1(x) \rho_{\text{ON}}(x, t; \tau) dx dt,$$

and define the relative *search efficiency* as

$$U(\tau, T) = \frac{\hat{U}(\tau, T)}{\hat{U}(0, T)}. \quad (19)$$

where $\hat{U}(0, T)$ is the uptake of the population of KS searchers. Therefore, $U(\tau, T) > 1$ indicates a gain in search efficiency, that is, increased stimulant encountered by the population due to the perceptual horizon $\tau > 0$.

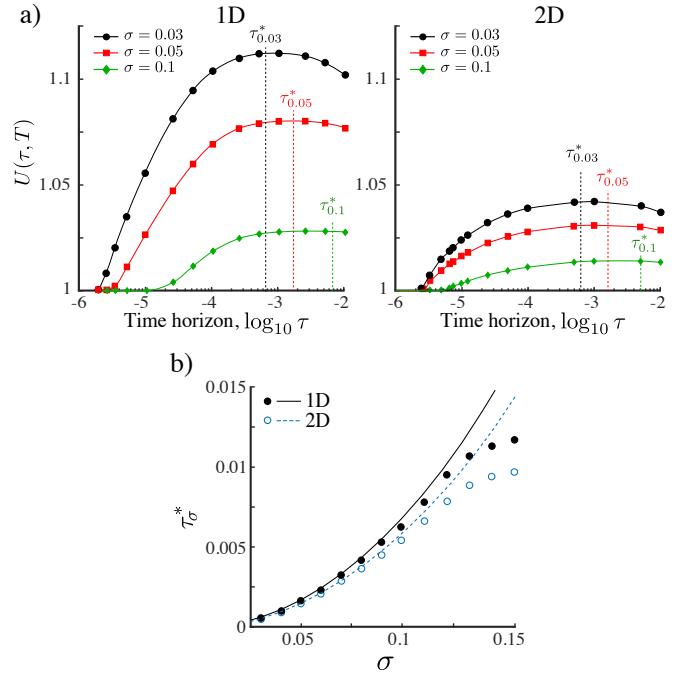


Figure 3. a) The relative search efficiency (19) of the ON population in one and two dimensions as a function of the time horizon τ for three Gaussian stimulant profiles with length scales $\sigma = 0.03, 0.05, 0.1$. All results are computed over a fixed search time $T = 10^{-1}$. Values of $U(\tau, T) > 1$ indicate improved search efficiency of the ON population as compared to the KS population. The maximum efficiency $U^*(\tau_\sigma^*, T)$ is achieved at a time horizon τ_σ^* . For a given Pe_1 , the relative search efficiency decreases in higher dimensions due to the increasing dominance of diffusive motion relative to ballistic motion. b) Comparison of simulations (circles) with our estimate (line) of τ_σ^* obtained from (20)–(21) by matching the mean-squared distance travelled to the length scale of the landscape as given by Eq. (22). The estimate is accurate when $\sigma \ll 1$.

Our numerics in Fig. 3a show that, given a finite search time T , the search efficiency (19) reaches a maximum $U^*(\tau_\sigma^*, T)$ for searchers operating with an optimal horizon τ_σ^* , which depends on the length scale σ for a given dimension d (Fig. 3a). The presence of a maximum follows from the asymptotic behaviour $U(\tau, T) \rightarrow 1$ for $\tau \rightarrow 0$ and $\tau \rightarrow \infty$. The latter limit follows from the invariance of $\rho_\infty(x)$ under τ , and the fact that the integral (19) is asymptotically dominated by the steady state. The presence of a maximum is also observed in the two-dimensional case, but $U^*(\tau_\sigma^*, T)$ decreases in higher dimensions.

The dependence of τ_σ^* with the length scale of the landscape σ obtained numerically from our simulations is shown in Fig. 3b with solid (1D) and open (2D) circles. To understand this dependence, consider a searcher at $x(t)$ obeying the Langevin equation (1) under the ON model. The reachable set until $t + \tau$ is within a ball of

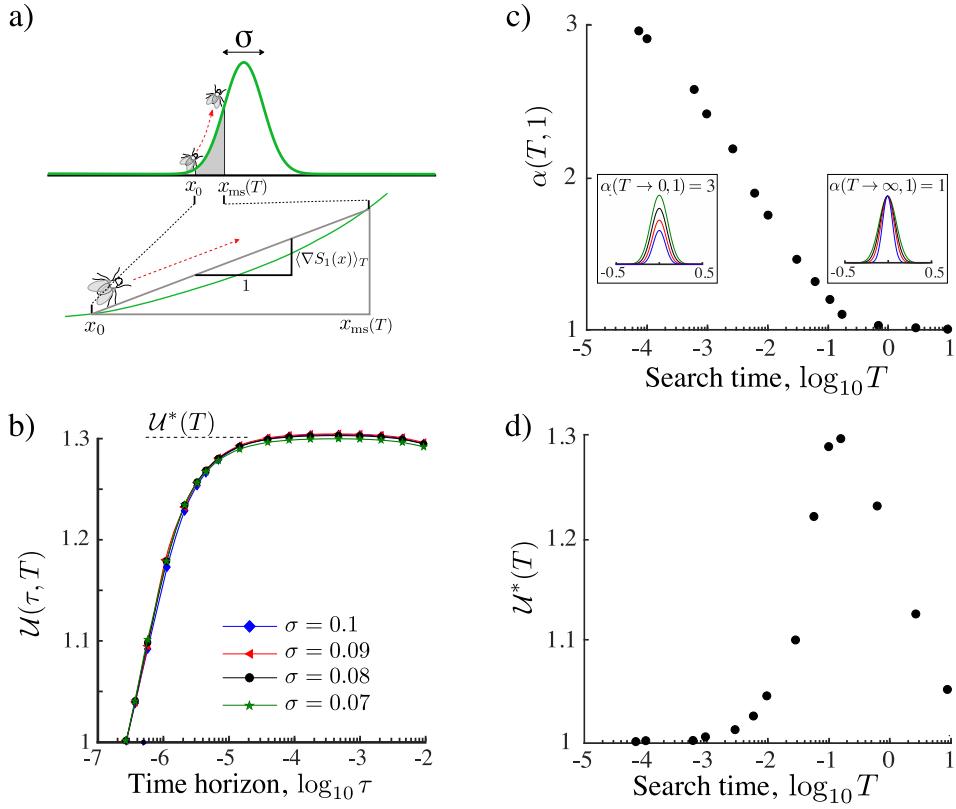


Figure 4. The search efficiency can be made invariant by scaling the sensitivity Pe_1 . a) The displacement of the searcher up to time T is approximated by the displacement in an effective gradient (23). To account for the change in this effective gradient as T or σ varies, we renormalise the Péclet number as in (25). b) The renormalisation yields invariance (26) of the search efficiency under σ ($T = 10^{-1}$ in this figure). c) The $\alpha(T, d)$ for $d = 1$ computed numerically matches the asymptotic results (24): $\alpha(T \rightarrow 0, 1) = 3$ and $\alpha(T \rightarrow \infty, 1) = 1$. From the perspective of the searcher, the rescaling $\tilde{Pe}_1(\sigma, T)$ is equivalent to renormalising the landscape. Insets show examples of the renormalised Gaussian landscape with varying σ but the same effective gradient. d) The optimal renormalised search efficiency $U^*(T)$ attainable by adapted agents achieves a maximum at search time $T^* \sim 0.1$. Hence non-local search is maximally advantageous for foraging times much smaller than the diffusion time $T_D = 1$.

radius $x_{ms}(\tau)$ defined by

$$x_{ms}^2(\tau) \simeq 2d\tau + \bar{v}^2(\tau)\tau^2, \quad (20)$$

where the two terms represent the displacement due to diffusion and to an effective drift velocity \bar{v} over time τ , respectively. To estimate $\bar{v}(\tau)$, we assume that the search time T is large enough such that individual searchers have explored the whole domain (e.g., for $T = 0.1$ this is fulfilled when $\tau > 10^{-3}$ as seen in Fig. 2b). The effective drift velocity can then be approximated by the velocity of an average searcher (over the domain) that maximises its gain up to time τ ,

$$\bar{v}(\tau) \simeq 2Pe_1 \int_{-1/2}^0 \frac{S_1(x + x_{ms}(\tau)) - S_1(x)}{x_{ms}(\tau)} dx, \quad (21)$$

where the integral is evaluated along a one-dimensional

cross-section of the Gaussian landscape

$$\begin{aligned} \bar{v}(\tau) &\simeq \frac{2Pe_1}{x_{ms}(\tau)} (2\pi\sigma)^{-\frac{d}{2}} \int_{-1/2}^0 \left(e^{-\frac{(x+x_{ms}(\tau))^2}{2\sigma^2}} - e^{-\frac{x^2}{2\sigma^2}} \right) dx \\ &= \frac{Pe_1}{x_{ms}(\tau)} (2\pi\sigma)^{-\frac{d-1}{2}} \left[\operatorname{erf} \left(\frac{x_{ms}(\tau)}{\sqrt{2\sigma^2}} \right) - \operatorname{erf} \left(\frac{1/2}{\sqrt{2\sigma^2}} \right) \right. \\ &\quad \left. - \operatorname{erf} \left(\frac{x_{ms}(\tau) - 1/2}{\sqrt{2\sigma^2}} \right) \right]. \end{aligned}$$

From (20) and (21), we obtain an estimate of the horizon τ necessary to search over a distance x_{ms} with the ON model.

The relevance of this estimate is shown in Fig. 3b, which shows that for small σ , the maximum search efficiency is attained when the mean-squared displacement of the searchers equals the length scale of the environment:

$$x_{ms}(\tau_\sigma^*) \simeq \sigma, \quad (22)$$

as obtained with our approximation. The ON model thus predicts that the most efficient searchers are those that

tune their horizon such that they traverse the characteristic length scale of the environment within one optimisation step. Shorter or longer optimisations lead to a decreased search efficiency.

The dependence of this behaviour on the dimension d is also captured by (20), which tells us that the ballistic and diffusive terms balance when $\bar{v}(\tau)^2\tau/(2d)$. Thus for a given S_1 , τ and Pe_1 the motion of searchers becomes gradually diffusion dominated as the dimension d increases. As a result, the relative search efficiency decreases in higher dimensions as shown in Fig. 3.

IV. INVARIANCE OF SEARCH EFFICIENCY THROUGH RESCALING OF RESPONSE SENSITIVITY

The search efficiency of the ON model depends on the length scale σ of the Gaussian landscape: the ON gain diminishes as σ increases (Fig. 3). However, as we now show, an ON searcher can retain the same search efficiency under a Gaussian landscape with varying length scale by adjusting Pe_1 , its sensitivity to the stimulant.

To see this, consider an ON searcher starting at x_0 exposed to its nutrient micro-environment until time T . The effective gradient for this searcher depends on the starting position x_0 and is given by:

$$\langle \nabla S_1(x) \rangle_T = \frac{\int_{x_0}^{x_0+x_{\text{ms}}(T)} \left[\frac{S_1(x+x_{\text{ms}}(\tau)) - S_1(x)}{x_{\text{ms}}(\tau)} \right] dx}{x_{\text{ms}}(T)}. \quad (23)$$

For a fixed exploration time T , an increase in the stimulant length scale σ leads to shallower effective gradients (Fig. 4a). Using asymptotic techniques, we show in Sect. S1 in the SI that the effective gradient (23) for symmetric Gaussian profiles has a well-defined behaviour in the two limiting regimes:

$$\begin{cases} \langle \nabla S_1(x) \rangle_T \propto \sigma^{-2-d} & \text{as } T \rightarrow 0, \\ \langle \nabla S_1(x) \rangle_T \propto \sigma^{-d} & \text{as } T \rightarrow \infty. \end{cases} \quad (24)$$

Together with the form of the dynamics (1), this suggests the following scaling for the Péclet number:

$$\widetilde{\text{Pe}}_1(\sigma, T) \propto \sigma^{\alpha(T, d)}, \quad (25)$$

with $d \leq \alpha(T, d) \leq 2 + d$.

We have tested this scaling by obtaining the ON solution (13) over a given T for Gaussian profiles with different σ using the renormalised Péclet number (25). We then compute the relative search efficiency (19) for this solution, $U(\tau, T; \widetilde{\text{Pe}}_1(\sigma, T))$. Our numerics in Fig. 4b show that the search efficiency curves for the renormalised parameter (25) for different σ collapse on a single curve:

$$U(\tau, T; \widetilde{\text{Pe}}_1(\sigma, T)) \approx \mathcal{U}(\tau, T). \quad (26)$$

The exponent $\alpha(T, 1)$, i.e. for $d = 1$, is obtained numerically (Fig. 4c) is consistent with the expected asymptotic limits (24). Note that, the effective gradient (23) is a function of $x_{\text{ms}}(\tau)$, which changes with dimension. Therefore, although for any d a scaling relationship exists, it will be different from the curve in Fig. 4c depending on d .

Hence the search efficiency can be made invariant for different environmental length scales σ by rescaling the Péclet number (25). Alternatively, adjusting $\alpha(T, d)$ can be viewed as responding to a 'renormalised landscape' [$\widetilde{\text{Pe}}_1 S_1(x)$] in order to maintain the ON search efficiency. This is intuitive in limiting cases: when the search time is small ($T \rightarrow 0$), the efficiency remains unchanged on landscapes with similar local gradients near the centre ($x_0 \ll 1$); when the search time is large ($T \rightarrow \infty$), the efficiency is invariant for landscapes with similar effective gradient over the whole domain (see inset of Fig. 4c).

This result suggests that searchers can optimise their search efficiency by adjusting their response sensitivity (as in the scaling (25)) so as to balance the relative effect of the advection and diffusion velocities or, in other words, the relative importance of gradient optimisation versus noisy exploration. Since the diffusion coefficient D is typically independent of $S_1(x)$ [27], the adjustment of Pe_1 could be achieved by varying the sensitivity as a function of the stimulant, i.e., $\chi_1(S_1(x))$ (see Appendix A). In the Discussion, we explain possible biological mechanisms to achieve this effect.

Is non-local search advantageous over search times relevant for ecological systems? The invariant search efficiency $\mathcal{U}(\tau, T)$ characterises the performance of a searcher that is tuned to the intrinsic length scale of the stimulant landscape during its search time T . In Fig. 4d, we show the dependence of the maximum renormalised search efficiency (26) with the search time T . As expected, for short search time $T \rightarrow 0$ and long search time $T \rightarrow \infty$, the efficiency $\mathcal{U}^*(T)$ is equivalent to the local search strategy, i.e., $\mathcal{U}^*(T) \rightarrow 1$. However, between both extremes, searchers benefit from finite perception. Our numerics show that the optimal search time is $T^* = \text{argmax } \mathcal{U}^*(T, \tau^*) \sim 0.1 < 1 = T_D$. Hence finite perception is maximally advantageous for search times smaller than the diffusion time, a fact that is typical in ecological systems (see Discussion).

V. AGENT INTERACTION IN THE OPTIMAL NAVIGATION MODEL: INCREASED EFFICIENCY AND MULTIMODALITY

Until now, we have considered a population of non-interacting searchers reacting only to an external stimulant $S_1(x)$. Now we consider searchers which also interact with each other through an attractive secondary stimulant $S_2(x, t)$ (i.e., $\text{Pe}_2 > 0$). The effect of such interaction can be potentially contradictory: agent interaction may increase aggregation; however, aggregation might not in-

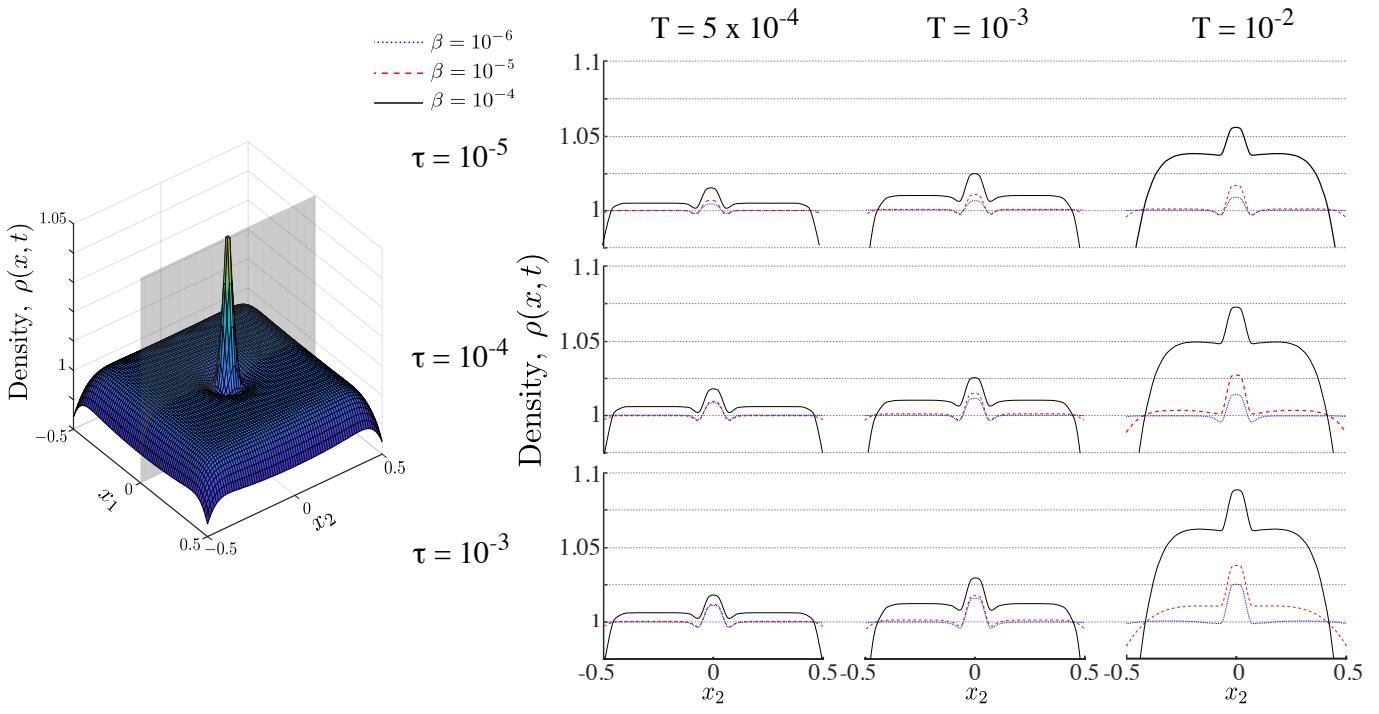


Figure 5. Populations of interacting ON agents. Right panel shows snapshots of a population of ON searchers along the cutting plane shown in grey on the left hand panel. Snapshots are shown with time horizons τ and interaction strength β for different search times T . Interaction increases aggregation at the maximum of the nutrient profile (the centre of the domain) and reduces the long tails leading to the emergence of multimodality. Simulations performed with $\text{Pe}_1 = 2$, $\sigma = 0.05$.

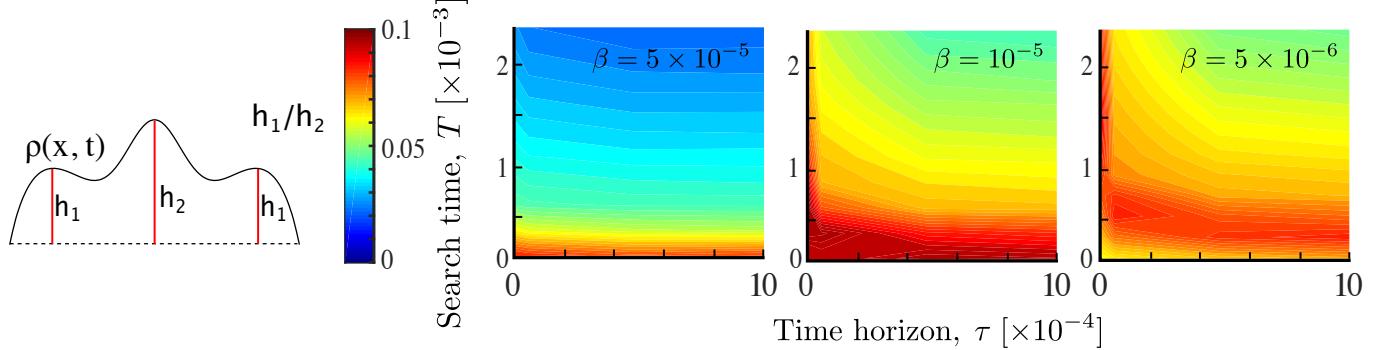


Figure 6. Multimodality as a result of agent interaction. The relative prominence of the smaller modes h_1/h_2 is computed for different parameters T , τ and β . As the interaction β decreases, multimodality gradually shifts to longer times and its magnitude also depends on the time horizon. This results from non-local searchers reaching more quickly a unimodal aggregate at the stimulant source. Simulations performed with $\text{Pe}_1 = 2$, $\sigma = 0.05$.

crease search efficiency for the population if a large proportion of the agents clumps away from the stimulant source if the sensitivity to stimulant S_2 becomes larger.

To explore these effects numerically, we use the ON model (13) with a weak interaction in the free energy (10)

$$\text{Pe}_2 = \beta \text{Pe}_1, \quad \beta \ll 1 \quad (27)$$

and compute the time evolution of the population for various β and τ . We restrict ourselves to the weak interaction case ($\beta \ll 1$) to facilitate our numerics. Specif-

ically, to prevent infinite density concentration at finite time, a well-known artefact of the KS model, by applying a small regularising factor $\omega \rho^2$ with $0 < \omega \ll 1$ to the free energy $\mathcal{F}[\rho]$ (10). This is a volume exclusion term that models the fact that agents occupy a finite volume in space [36, 37] (see (B5) in Appendix B). The weak interaction assumption ($\beta \ll 1$) is important for numerical performance so as to avoid high velocities near the concentration point that would necessitate a finer time discretisation to maintain the CFL condition (see Appendix B).

A summary of the numerics of the two-dimensional ON case with different horizons τ and levels of interaction β over different search times T is presented in Fig. 5. In all cases, the presence of interaction reduces the tails of the population density and increases aggregation near the maximum of S_1 at the centre of the domain. This becomes more noticeable as the search time grows and we approach stationarity (Fig. 5, right column). During transients, however, agent interaction induces multimodality in the population (Fig. 5, left and middle column). This implies that some searchers move away from the maximum of the stimulant $S_1(x)$ and aggregate into transient subpopulations. This behaviour arises due to the non-linear response of the searchers to the gradient of $S_1(x)$: for steep gradients of $S_1(x)$, agents are driven by attraction to S_1 , whereas for shallow gradients of $S_1(x)$ the agents are driven by the interaction with the rest of the population through the secreted stimulant S_2 . In contrast, for non-interacting agents ($\beta = 0$) the distribution is always uni-modal (Fig. 5).

Fig. 6 shows the dependence of multimodality on T and τ at three different interaction strengths $\beta = 5 \times 10^{-5}, 10^{-5}, 5 \times 10^{-6}$. We find that interaction introduces multimodality, as part of the population clumps away from the source of the attractant. For higher interaction strength ($\beta = 5 \times 10^{-5}$) multimodality is only present during early transients, and is independent of the time horizon τ , since the strongly interacting population rapidly converges to the peak of the primary stimulant. As the interaction strength β decreases, multimodality lasts longer, but it decreases as the time horizon τ increases. This is in line with the expectation that increased spatial perception τ or search time T leads to overlapping information of the searchers about the environment, whereas, in contrast, when τ and/or T are small, the searchers remain isolated within their local environment, thus leading to multimodality by local aggregation. As the interaction decreases further multimodality disappears.

Importantly, despite the presence of such mildly multimodal transients with clumped subpopulations, agent interaction β leads to an overall improvement in the search efficiency $U^*(\tau^*, T)$ at a shorter time horizon τ^* (Fig. 7). This behaviour follows from the increased concentration of searchers around the centre (on average), and suggest that when resources are sparsely distributed, agent interaction in conjunction with a finite perceptual range, could play a role in improving the collective sensing of the environment towards improved population efficiency.

VI. DISCUSSION

In this work, we studied how finite perception influences the dynamics and efficiency of collective random search in a population of agents. Using concepts from optimal control and random walks, we proposed a model that encapsulates the spatial information the searchers

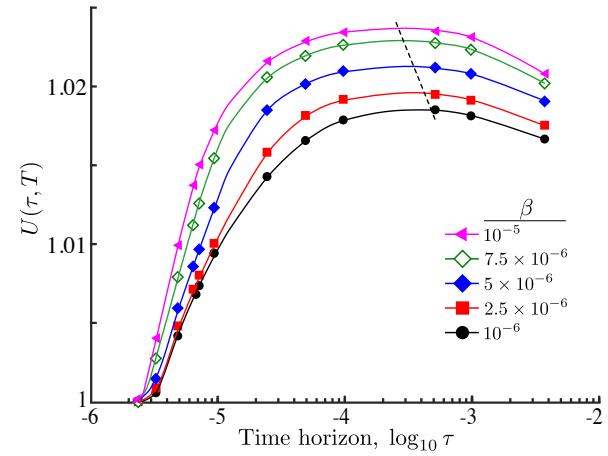


Figure 7. Relative search efficiency (19) of an ON population of interacting agents against time horizon with different interaction strengths β . As β grows, the search efficiency increases and the maximum is attained at smaller time horizons (dashed line). The improvement in the search efficiency with interaction follows from the increased accumulation of agents at the maximum of the nutrient and the reduced long tails far from the centre of the domain (Fig. 5). Simulation parameters: $Pe_1 = 2$, $\sigma = 0.05$, $T = 5 \times 10^{-2}$.

possess as a time horizon for an optimisation problem. Simulations of the dynamics of population search show that non-local information affects the movement strategy, as compared with the standard Keller-Segel model based on local optimisation. Although non-local search does not change the stationary state, it leads to qualitatively different transient responses of possible relevance in biological systems [7]. For example, marine bacteria have been observed to aggregate at food patches much faster ($\sim 10^1 - 10^2$ s) [38, 39] than the timescale to reach the steady state distribution ($T_D \sim 10^4 - 10^5$ s, based on $D \sim 10^2 - 10^3 \mu\text{m}^2/\text{s}$ [40] and typical inter-patch distance $L \sim 10^3 \mu\text{m}$ [41]). Similarly, in several rodent species diffusion coefficients of $D \sim 200 \text{ m}^2/\text{day}$ and home range of $L \sim 70 \text{ m}$ have been reported [42]. Therefore, the time it takes to reach home by diffusion (~ 25 days) is much longer than their typical response time (~ 1 day).

We find that non-local information can account for the increase in the search efficiency under transient search times: the maximum efficiency is reached when the mean-squared displacement of the searchers matches the environmental length scale of the stimulant. When the time horizon vanishes or when the search time is infinite, our model recovers the response of local searchers. This is in accordance with the fact that when long-range cues are unreliable, local response leads to highest efficiency [15].

We also showed that the search efficiency can be made invariant to changes in environmental length scales by suitably scaling the response sensitivity. As a consequence, a searcher with a given perceptual range may always achieve its maximum efficiency by dynamically adjusting its sensitivity to the environmental stimuli.

This can be achieved either by dynamically rescaling the responses via adaptation at the agent level [43], or by the presence of a distribution of sensitivities among the agents at the population level. For example, it has been shown that phenotypic heterogeneity (or plasticity) across a population can be used to achieve maximum search efficiency in patchy environments [44].

Finally, we considered the effect of interaction between searchers with finite perception and showed that interaction can lead to unimodal or multimodal population distributions on transient timescales. Multimodality appears even in the presence of unimodal stimulant landscapes due to a trade-off between following the environmental gradient or the rest of the population. In our numerics, interaction always improved the overall search efficiency of the population.

Recent theoretical and experimental studies [45, 46] also suggest that rodents (and other higher animals) store spatial information of their environment, and in doing so commonly prefer longer (non-trivial) future paths as opposed to paths leading to immediate rewards. These results are in line with a predictive optimisation based on spatial knowledge, as in the ON model, where the searcher weighs up local cues with those at a distance to inform the planning of the future trajectory as opposed to an immediate (gradient) optimisation. Clearly, the perceptual horizon in the ON model is a very coarse abstraction of the spatial information the organism perceives (for example in its visual field) or is encoded in its neural or cellular memory. Going beyond our framework would require explicit assumptions about the sensory or cognitive mechanisms involved in search and navigation, an area of interesting future research which is out of the scope of this work.

Our work opens up several directions of research. Beyond our simple setup, it would be of interest to study search on temporally-fluctuating or patchy nutrient landscapes [47] using non-local strategies. Random search theory based on local response predicts that when a searcher is positioned equally far from two nutrient patches it is equally likely to explore either patch. However, on transient timescales, nonlocal searchers are expected to explore the patch with denser resources with higher probability. This direction will be the object of future work.

Appendix A: The Keller-Segel model and its nondimensionalisation

On space and time scales much larger than the microscopic motion of the agents, we describe the biased random walk of the searchers with a Langevin equation (1)-(2), and the time evolution of the population density, $\rho(x, t)$, with the associated Fokker-Planck equation [28]:

$$\partial_t \rho - D \nabla^2 \rho + \nabla \cdot (\rho \nabla (\chi_1 S_1 + \chi_2 S_2)) = 0. \quad (\text{A1})$$

If the agents release the stimulant S_2 at a constant rate Γ (in units of mol/(agent t)) and S_2 diffuses with diffusion coefficient D_2 , we obtain the forced diffusion equation:

$$\partial_t S_2 - D_2 \nabla^2 S_2 = \Gamma \rho. \quad (\text{A2})$$

To reduce the number of parameters, we nondimensionalise the equations. We use the length scale L and time scale $T_D = L^2/D$ to define

$$\hat{x} := x/L \quad \hat{t} := t/T_D.$$

Let $\hat{S}_2 := S_2/S_{2,\text{av}}$ and $\hat{\rho} := \rho/\rho_{\text{av}}$ with $\rho_{\text{av}} = \frac{1}{L^d} \int_{\Omega} \rho(x) dx$ and $S_{2,\text{av}} = \Gamma L^2 \rho_{\text{av}} / D_2$. Then Eq. (A2) can be rewritten in non-dimensional form:

$$\frac{D}{D_2} \partial_{\hat{t}} \hat{S}_2 - \nabla_{\hat{x}}^2 \hat{S}_2 = \hat{\rho}. \quad (\text{A3})$$

If S_2 diffuses faster than the searchers ($D/D_2 \ll 1$), the time-dependence in (A3) may be neglected to first order and the resulting Poisson equation has the quasi-steady solution:

$$\hat{S}_2(\hat{x}) = \int_{\Omega} \Phi(\hat{x} - x') \hat{\rho}(x', t) dx' := \Phi * \hat{\rho}, \quad (\text{A4})$$

where $\Phi(x)$, given by (4), is the Green's function solution to the Poisson equation in d dimensions.

Defining $\hat{S}_1 := S_1/S_{1,\text{av}}$ with $S_{1,\text{av}} = \frac{1}{L^d} \int_{\Omega} S_1(x) dx$ we rewrite Eq. (A1) in non-dimensional form:

$$\partial_{\hat{t}} \hat{\rho} - \nabla_{\hat{x}}^2 \hat{\rho} + \nabla_{\hat{x}} \cdot \left(\hat{\rho} \nabla_{\hat{x}} \left(\text{Pe}_1 \hat{S}_1 + \text{Pe}_2 \hat{S}_2 \right) \right) = 0, \quad (\text{A5})$$

where $\text{Pe}_1 = \chi_1 S_{1,\text{av}} / D$ and $\text{Pe}_2 = \chi_2 S_{2,\text{av}} / D$ are Péclet numbers. To simplify the notation, we drop the hats in all variables in the main body of the paper and always consider non-dimensionalised equations. From (A5) and (A4), we get the non-dimensional Keller-Segel model (5).

Appendix B: Algorithm and numerical scheme for the solution of the Optimal Navigation model

For notational simplicity we present the numerical solution of the ON model in the one-dimensional case. The extension to two dimensions is straightforward. The ON model (13) is numerically solved as an implicit forward Euler scheme, where we obtain during each step the evolution $m(x, s) = \rho(x, k\tau + s)$ for the time window $s \in [0, \tau]$ using a recently developed gradient descent algorithm [35]. We used this algorithm since it preserves a discrete analogue of the variational structure of the ON model and has the required property that the free energy (10) decays along the time evolution of $m(x, s)$.

ALGORITHM 1. Gradient descent algorithm for (13) based on Ref. [35]. Simulation parameters used: $\Delta x = 0.01$, $\Delta t = 10^{-6}$, $\kappa = 1$, $\delta = 10^{-6}$.

Input $\rho(x, k\tau)$, κ , δ

Output $m(x, s) := \rho(x, k\tau + s)$, $s \in [0, \tau]$,

Set $m_0^0 \leftarrow \rho(x, k\tau)$

while $\|u^{p+1} - u^p\|_{L^2} > \delta$ **do**

 Compute the FP equation (B3a) forward in time

for $j = 0 : 1 : N$ **do**

 Compute $m_{j+1}^{p+1} = (I - B(u_j^p))m_j^p$

 Set $u^{p+1} \leftarrow m_N^{p+1}$ and $\phi_N^{p+1} \leftarrow -\partial \mathcal{J}_{\Delta t, \Delta x}(u^{p+1})$

end for

 Compute the HJ equation (B3b) backward in time

for $j = N + 1 : -1 : 1$ **do**

 Compute $\phi_j^{p+1} = (I - B^T(u_j^p))\phi_{j+1}^p - \frac{\Delta t}{2}r_j^p$

end for

 Check $[\mathcal{J}_{\Delta t, \Delta x}(m^p, u^p) - \mathcal{J}_{\Delta t, \Delta x}(m^{p+1}, u^p)] < 0$

 Update $(1 + \tau)u^{p+1} = \nabla \phi^{p+1} + \kappa u^p$

 Set $p \leftarrow p + 1$

end while

Return m

We first derive the optimality conditions of (13). The corresponding Lagrangian is [33]:

$$\begin{aligned} \mathcal{L}^\tau(m, u, \phi) := & \mathcal{J}(u, m) \\ & + \int_0^\tau \left[\mathcal{F}[m(x, \tau)] - \int_\Omega (\partial_s \phi + u \cdot \nabla \phi) m \, dx \right] ds \quad (\text{B1}) \\ & + \int_\Omega [\phi(x, \tau)m(x, \tau) - \phi(x, 0)] m(x, 0) dx, \end{aligned}$$

allowing us to write (13) as a saddle node problem:

$$\inf_{(m, u)} \sup_{\phi} \mathcal{L}^\tau(m, u, \phi), \quad (\text{B2})$$

where the infimum is taken over all the possible time evolutions of the population, i.e. pairs (m, u) satisfying $\partial_s m + \nabla \cdot (mu) = 0$. Solving for the first-order optimality conditions of (B2),

$$\partial_m \mathcal{L}^\tau = 0, \quad \partial_u \mathcal{L}^\tau = 0, \quad \partial_\phi \mathcal{L}^\tau = 0,$$

yields a variational mean-field game system [48, 49] consisting of the Fokker-Planck equation (B3a) solved forward in time coupled to a Hamilton-Jacobi (HJ) equation (B3b) solved backward in time:

$$\partial_s m + \nabla \cdot (mu) = 0, \quad m(\cdot, 0) = \rho(\cdot, k\tau) \quad (\text{B3a})$$

$$\partial_s \phi + u \cdot \nabla \phi + \frac{|u|^2}{2} = 0, \quad \phi(\cdot, \tau) = -\frac{\delta \mathcal{F}[m]}{\delta m} \Big|_{s=\tau}, \quad (\text{B3b})$$

with $u = -\nabla \phi$. The terminal condition in (B3b) can be explicitly solved as

$$\frac{\delta \mathcal{F}[m]}{\delta m} = \log m(x, t) - \text{Pe}_1 S(x) + \text{Pe}_2 \log |x| * m.$$

Consider a uniform discretisation of $\Omega = [-1/2, 1/2] \times [0, \tau]$ as a grid (i, j) , $i = 0, \dots, M + 1$, $j = 0, \dots, N + 1$ with steps $\Delta x = 1/(M + 1)$ in space and $\Delta t = \tau/(N + 1)$ in time for positive integers M, N , where the first and last points are boundary nodes. Let $m_{i,j} := m(i\Delta x, j\Delta t)$ on the grid points and $u_{i+1/2,j} := u((i+1/2)\Delta x, j\Delta t)$ on the staggered grid points. Finally, we impose Neumann conditions at the boundaries, $u_{1/2} = u_{M+3/2} = 0$, and denote the array of point excluding the two boundary nodes by $u_j = (u_{i+1/2,j})_{i=1, \dots, M}$ and $m_j = (m_{i,j})_{i=1, \dots, M}$.

We discretise the cost function $\mathcal{J}[m, u]$ as

$$\mathcal{J}_{\Delta t, \Delta x}(m, u) := \Delta x \Delta t \sum_{i=1}^M \sum_{j=0}^N \left(\frac{1}{2} r_{i,j} m_{i,j} + f_i(m_N) m_{i,N} \right), \quad (\text{B4})$$

where $r_{i,j} = ((u_{i-1/2,j})^2 + (u_{i+1/2,j})^2)/2$,

$$\begin{aligned} f_i(m_j) = & \log(m_i) - \text{Pe}_1 S + \text{Pe}_2 \sum_{l=1}^M \log((|i-l|\Delta x) + \epsilon) m_{i,j} \\ & + \omega m_{i,j}^2. \end{aligned} \quad (\text{B5})$$

Here $\epsilon = 10^{-5}$ and $\omega = 10^{-6}$ are small regularising factors to bound the logarithmic term and to guarantee the existence of finite solutions at all times [36, 37].

We use a first order upwind scheme for the advection in (B3a):

$$m_{i,j+1} = m_{i,j} + \frac{\Delta t}{\Delta x} (m_{i+1/2,j} u_{i+1/2,j} - m_{i-1/2,j} u_{i-1/2,j}),$$

$$\text{where } m_{i+1/2,j} = \begin{cases} m_{i+1,j}, & \text{if } u_{i+1/2,j} < 0 \\ m_{i,j}, & \text{if } u_{i+1/2,j} \geq 0. \end{cases}$$

$$u_{1/2,j} = u_{M+1/2,j} = 0, \quad j = 0, \dots, N \quad (\text{Neumann BC's}).$$

In compact form this may be written as

$$m_{j+1} = (I - B(u_j))m_j, \quad (\text{B6})$$

where $I \in \mathbb{R}^{M \times M}$ is the identity matrix and $B(u) \in \mathbb{R}^{M \times M}$ is the matrix associated with the upwind scheme.

The corresponding scheme for (B3b) is given by

$$\phi_j = (I - B^T(u_j))\phi_{j+1} + \frac{\Delta t}{2} r_j. \quad (\text{B7})$$

Algorithm (1) may then be used with a given descent step κ to compute $m(x, s)$ to the required tolerance level δ . Note that κ , Δt and Δx must be chosen so as to fulfil the CFL condition [50]: $u_{i+1/2,j} \Delta t / \Delta x \leq 1$ for all staggered grid points $(i+1/2, j)$. Then the cost function $\mathcal{J}_{\Delta t, \Delta x}$ decreases at every iteration p .

ACKNOWLEDGEMENTS

AG acknowledges funding through a PhD studentship under the BBSRC DTP at Imperial College (BB/M011178/1). JAC acknowledges funding from the EPSRC EP/P031587/1. MB acknowledges funding from the EPSRC (EP/N014529/1 and EP/I032223/1).

[1] A. M. Hein, F. Carrara, D. R. Brumley, R. Stocker, and S. A. Levin, “Natural search algorithms as a bridge between organisms, evolution, and ecology,” *Proc. Natl. Acad. Sci. U.S.A.* **113**, 9413–9420 (2016).

[2] W. J. Bell, *Searching Behaviour*, The behavioural ecology of finding resources (Springer Netherlands, 2012).

[3] G. M. Viswanathan, M. G. E. da Luz, E. P. Raposo, and H. E. Stanley, *The Physics of Foraging*, An Introduction to Random Searches and Biological Encounters (Cambridge University Press, 2011).

[4] D. W. Stephens and J. R. Krebs, *Foraging theory* (Princeton University Press, 1986).

[5] R. H. MacArthur and E. R. Pianka, “On optimal use of a patchy environment,” *Am. Nat.* **100**, 603–609 (1966).

[6] E. L. Charnov, “Optimal foraging, the marginal value theorem,” *Theor. Popul. Biol.* **9**, 129 – 136 (1976).

[7] A. Hastings, “Transients: the key to long-term ecological understanding?” *Trends. Ecol. Evol.* **19**, 39–45 (2003).

[8] N. Strelkowa and M. Barahona, “Switchable genetic oscillator operating in quasi-stable mode,” *J. R. Soc. Interface*, rsif20090487 (2010).

[9] N. Strelkowa and M. Barahona, “Transient dynamics around unstable periodic orbits in the generalized repressor model,” *Chaos* **21**, 023104 (2011).

[10] E. Schrödinger, *What is life?: With mind and matter and autobiographical sketches* (Cambridge University Press, 1992).

[11] H. C. Berg, *Random walks in biology* (Princeton University Press, 1993).

[12] O. Bénichou, M. Coppey, M. Moreau, P-H. Suet, and R. Voituriez, “Optimal search strategies for hidden targets,” *Phys. Rev. Lett.* **94**, 198101 (2005).

[13] V. Tejedor, R. Voituriez, and O. Bénichou, “Optimizing persistent random searches,” *Phys. Rev. Lett.* **108**, 088103 (2012).

[14] A. Okubo and S. A. Levin, *Diffusion and Ecological Problems: Modern Perspectives* (Springer-Verlag New York, 2001).

[15] M. Vergassola, E. Villermaux, and B. I. Shraiman, “Infotaxis’ as a strategy for searching without gradients,” *Nature* **445**, 406–409 (2007).

[16] F. Bartumeus and J. Catalan, “Optimal search behavior and classic foraging theory,” *J. Phys. A Math. Theor.* **42**, 434002 (2009).

[17] 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**, 911 EP –.

[18] E. I. Moser, E. Kropff, and M-B Moser, “Place cells, grid cells, and the brain’s spatial representation system,” *Annu. Rev. Neurosci.* **31** (2008).

[19] D. C. Mistro, L. A. D. Rodrigues, and W. C. Ferreira, “The africanized honey bee dispersal: a mathematical zoom,” *Bull. Math. Biol.* **67**, 281 – 312 (2005).

[20] A. Gomez-Marin, G. J. Stephens, and M. Louis, “Active sampling and decision making in *Drosophila* chemotaxis,” *Nat. Commun.* **2**, 441 (2011).

[21] A. Mitchell, G. H. Romano, B. Groisman, A. Yona, E. Dekel, M. Kupiec, O. Dahan, and Y. Pilpel, “Adaptive prediction of environmental changes by microorganisms,” *Nature* **460**, 220–224 (2009).

[22] R. van den Bos, “Animal anticipation: A perspective,” in *Handbook of Anticipation: Theoretical and Applied Aspects of the Use of Future in Decision Making*, edited by R. Poli (Springer, 2017) pp. 1–13.

[23] A. Gosztolai, J. Schumacher, V. Behrends, J. G. Bundy, F. Heydenreich, M. H. Bennett, M. Buck, and M. Barahona, “Glnk facilitates the dynamic regulation of bacterial nitrogen assimilation,” *Biophys. J.* **112**, 2219–2230 (2017).

[24] E. A. Codling, M. J. Plank, and S. Benhamou, “Random walk models in biology,” *J. R. Soc. Interface* **5**, 813–834 (2008).

[25] E. F. Keller and L. A. Segel, “Initiation of slime mold aggregation viewed as an instability,” *J. Theor. Biol.* **26**, 399–415 (1970).

[26] D. Boyer, D. S. Dean, C. Mejía-Monasterio, and G. Oshanin, “Optimal estimates of the diffusion coefficient of a single brownian trajectory,” *Phys. Rev. E* **85**, 031136 (2012).

[27] H. G. Othmer, S. R. Dunbar, and W. Alt, “Models of dispersal in biological systems,” *J. Math. Biol.* **26**, 263–298 (1988).

[28] C. Gardiner, *Stochastic methods* (Springer, 2009).

[29] J. A. Carrillo, A. Chertock, and Y. Huang, “A finite-volume method for nonlinear nonlocal equations with a gradient flow structure,” *Commun. Comput. Phys.* **17**, 233–258 (2015).

[30] J. A. Carrillo, R. J. McCann, and C. Villani, “Kinetic equilibration rates for granular media and related equations: entropy dissipation and mass transportation estimates,” *Rev. Mat. Iberoamericana* **19**, 971–1018 (2003).

[31] R. Jordan, D. Kinderlehrer, and F. Otto, “The variational formulation of the Fokker-Planck equation,” *SIAM J. Math. Anal.* **29**, 1–17 (1998).

[32] A. Blanchet, V. Calvez, and J. A. Carrillo, “Convergence of the mass-transport steepest descent scheme for the subcritical Patlak-Keller-Segel model,” *SIAM J. Numer. Anal.* **46**, 691–721 (2008).

[33] J.-D. Benamou and Y. Brenier, “A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem,” *Numer. Math.* **84**, 375–393 (2000).

[34] J. A. Carrillo and J. S. Moll, “Numerical simulation of diffusive and aggregation phenomena in nonlinear continuity equations by evolving diffeomorphisms,” *SIAM J.*

Sci. Comput. **31**, 4305–4329 (2009/10).

- [35] M. Burger, M. Franek, and C. B. Schönlieb, “Regularized regression and density estimation based on optimal transport,” Appl. Math. Res. Express. **2012**, 209–253 (2012).
- [36] V. Calvez and J. A. Carrillo, “Volume effects in the Keller-Segel model: energy estimates preventing blow-up,” J. Math. Pures Appl. (9) **86**, 155–175 (2006).
- [37] J. A. Carrillo, S. Hittmeir, and A. Jägel, “Cross diffusion and nonlinear diffusion preventing blow up in the Keller-Segel model,” Math. Models. Methods. Appl. Sci. **22**, 1250041 (2012).
- [38] N. Blackburn, T. Fenchel, and J. Mitchell, “Microscale nutrient patches in planktonic habitats shown by chemotactic bacteria,” Science **282**, 2254–2256 (1998).
- [39] J. R. Seymour, R. Simó, T. Ahmed, and R. Stocker, “Chemoattraction to dimethylsulfonylpropionate throughout the marine microbial food web,” Science **329**, 342–345 (2010).
- [40] T. Kiørboe, H.-P. Grossart, H. Ploug, and K. Tang, “Mechanisms and rates of bacterial colonization of sinking aggregates,” Appl. Environ. Microbiol. **68**, 3996–4006 (2002).
- [41] F. Azam and J. W. Ammerman, “Cycling of organic matter by bacterioplankton in pelagic marine ecosystems: microenvironmental considerations.” in *Marine Phytoplankton and Productivity*, edited by O. Holm-Hansen, L. Bolis, and R. Gilles (Springer, Berlin, 1984).
- [42] L. Giuggioli, G. Abramson, V. M. Kenkre, G. Suzán, E. Marcé, and T. L. Yates, “Diffusion and home range parameters from rodent population measurements in panama,” Bull. Math. Biol. **67**, 1135 (2005).
- [43] M. D. Lazova, T. Ahmed, D. Bellomo, R. Stocker, and T. S. Shimizu, “Response rescaling in bacterial chemotaxis,” Proc. Natl. Acad. Sci. U.S.A. **108**, 13870–13875 (2011).
- [44] J. Hollander, M. Verzijden, E. Svensson, and C. Brönmark, “Dispersal and phenotypic plasticity,” in *Animal Movement Across Scales*, edited by L-A. Hansson and S. Akesson (Oxford University Press, 2014) pp. 110–125.
- [45] M. G. Mattar and N. D. Daw, “Prioritized memory access explains planning and hippocampal replay.” Nat. Neurosci. **21**, 1609–1617 (2018).
- [46] B. E. Pfeiffer and D. J. Foster, “Hippocampal place-cell sequences depict future paths to remembered goals.” Nature **497**, 74–79 (2013).
- [47] M. Chupeau, O. Bénichou, and S. Redner, “Search in patchy media: Exploitation-exploration tradeoff,” Phys. Rev. E **95**, 012157 (2017).
- [48] J. M. Lasry and P. L. Lions, “Mean field games,” Japanese Journal of Mathematics (2007).
- [49] M. Huang, R. P. Malhamé, and P. E. Caines, “Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle,” Commun. Inf. Syst. **6**, 221–252 (2006).
- [50] R. J. LeVeque, *Numerical methods for conservation laws* (Springer Science & Business Media, 1992).