Featured in Physics

Collective olfactory search in a turbulent environment

Mihir Durve,[1,2] Lorenzo Piro,[3,4] Massimo Cencini,[5] Luca Biferale,[3] and Antonio Celani[2]

[1]*Department of Physics, Università degli Studi di Trieste, Trieste 34127, Italy*

[2]*Quantitative Life Sciences, The Abdus Salam International Centre for Theoretical Physics–ICTP, Trieste 34151, Italy*

[3]*Department of Physics and INFN, University of Rome Tor Vergata, Via della Ricerca Scientifica 1, 00133 Rome, Italy*

[4]*Max Planck Institute for Dynamics and Self-Organization, Am Fassberg 17, 37077 Goettingen, Germany*

[5]*Istituto dei Sistemi Complessi, CNR, via dei Taurini 19, 00185 Roma, Italy*

(Received 22 November 2019; revised 6 March 2020; accepted 3 June 2020; published 7 July 2020)

Finding the source of an odor dispersed by a turbulent flow is a vital task for many organisms. When many individuals concurrently perform the same olfactory search task, sharing information about other members' decisions can potentially boost the performance. But how much of this information is actually exploitable for the collective task? Here we show, in a model of a swarm of agents inspired by moth behavior, that there is an optimal way to blend the private information about odor and wind detections with the public information about other agents' heading direction. Our results suggest an efficient multiagent olfactory search algorithm that could prove useful in robotics, e.g., in the identification of sources of harmful volatile compounds.

DOI: 10.1103/PhysRevE.102.012402

I. INTRODUCTION

Animals are often on the move to search for something: a food source, a potential mate, or a nesting site. In many instances their navigation is informed by airborne chemical cues. One of the best known, and most impressive, olfactory search behavior is displayed by male moths [1–4]. Males are attracted by pheromones emitted in minute amounts by females that might be hundreds of meters away. Air turbulence breaks the odor plume into small, sparse patches interspersed by clean air or other odors making conventional search strategies like gradient climbing highly ineffective [5,6]. Experimental studies have in fact shown that male moths display a different search strategy. They alternate between two phases, depending on whether a pheromone signal has been detected or not: surging, i.e., sustained upwind flight, and casting, i.e., extended alternating crosswind motion. This strategy and others have inspired the design of robotic systems for locating sources of harmful volatile compounds [7–11]. Albeit the effectiveness of individual search is already remarkable in itself, the performance can be further boosted by cooperation among individuals, even in the absence of a centralized control [12–18].

In this paper we tackle the problem of collective olfactory search in turbulent environments. When the search takes place in a group, two classes of informative cues are available to the agents. First, there is private information: the detection of external signals—odor, wind velocity, etc.—by an individual. This perception takes place at short distances and is not shared with group members. Second, there is public information (social cues) in the form of the decisions made by other individuals, accessible to (a subset of) other peers, usually by visual cues, and therefore with a longer transmission range. Since the action taken by another individual may be also informed by its own private perception of external inputs, public cues indirectly convey information about odor distribution and wind direction at a distance. However, the spatial and temporal filtering induced by the sharing of public cues may in principle destroy the relevant, hidden information about the external guiding signals.

These considerations naturally lead to the question of whether the public information is exploitable at all for the collective search process. And if it is, how should the agents combine private and public information to improve the search performances ? Below, we will address these questions by exploiting a combination of models for individual olfactory search and flocking behavior in a turbulent flow with a constant mean wind.

The article is organized as follows. In Sec. II we describe the model for the collective search in a flow. In Secs. III and IV we present the results of the multiagent search in a stochastic flow and in a more realistic one obtained by direct numerical simulations of the Navier-Stokes equations. Section V is devoted to conclusions and discussion. Appendix A details the algorithmic implementation of the single-agent search strategy. Appendix B describes the flow characteristics. Details on the parameters of the search algorithm and of the flows are presented in Appendix C. Finally, a systematic study of the robustness of the results upon changing some of the main parameters of the model and of the flow is presented in Appendix D.

II. A MODEL FOR COLLECTIVE OLFACTORY SEARCH

The model setup is illustrated in Fig. 1(a). Initially, N agents are randomly placed within a circle of radius R_b at a distance L_x from an odor source S, which emits J odor particles per unit time. The odor particles, representing patches of odor with a concentration above the detection threshold, are transported in the surrounding environment by a turbulent

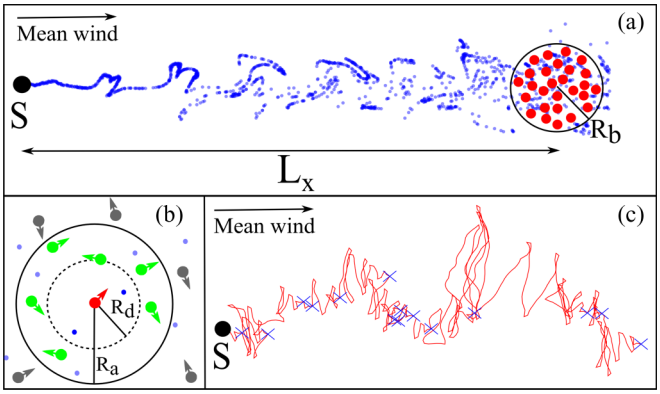

FIG. 1. Collective olfactory search. (a) Odor particles (blue dots) are emitted by the source S and dispersed by the turbulent flow. Agents (red) are initially placed far from S in a packed configuration. (b) Perception of an agent (red). Detected odor particles by the agent are shown as darker blue dots and agent's neighbors in green. Arrows indicate the instantaneous moving direction of agents. We set $L_x = 250R_d$, $R_b = 25R_d$, $R_a = 5R_d$, $R_d = 0.2$, $b = 2.5$. (c) Trajectory of an isolated agent performing the cast-and-surge program (see text). Blue crosses mark the detection of odor particles.

flow $\boldsymbol{u}$ with mean wind $\boldsymbol{U}$ (detailed below). The entire system is placed inside a larger square box of size bL_x with reflecting boundary conditions for the agents. A list of parameters with their numerical values is given in Appendix C.

We remark that in our numerical simulations we always placed the initial cloud of agents downwind of the source and centered along the mean-wind axis, as shown in Fig. 1. If the initial location is displaced away from the mean-wind axis, the search time dramatically increases as a consequence of the reduction of odor encounters (see, e.g., [19] about the case of a single agent in a diffusive environment). Since here we are focusing on assessing the role of public cues in the search strategy, we did not systematically study the performance at varying the initial crosswind location of the group of agents. Such an investigation, quite demanding from the computational viewpoint, could be the subject of further research.

In the next two sections, Secs. II A and II B, we separately describe the search strategies based either only on private cues or only on public cues. In the following section, Sec. II C, we then proceed to show how these schemes can be combined into a single algorithm blending both types of information.

A. Response to private cues

The behavior elicited by private cues such as odor and wind speed is inspired by the cast-and-surge strategy observed in moths. We adopted a modified version of the "active search model" [19] that works as follows. We assume that the agents have access to an estimate of the mean velocity of the wind, as moths actually do via a mechanism named optomotor anemotaxis [20]. In the model this estimate $\hat{\boldsymbol{u}}(t)$ is an exponentially discounted, running average of the flow velocity $\boldsymbol{u}$ perceived by the agent along its trajectory:

$$\hat{\boldsymbol{u}}(t) = \lambda \int_0^t \boldsymbol{u}(s) \exp[-\lambda(t-s)] ds. \tag{1}$$

The parameter λ is the inverse of the memory time: for $\lambda \to 0$ the estimate converges to the mean wind, while for $\lambda \to \infty$ it reduces to the instantaneous wind velocity at the agent current location. In the following we set $\lambda = 1$, which is of the same order of magnitude of the inverse correlation time of the flow. Notice that the estimated wind provides contextual information only about the location of the source. Indeed, in our model the agents are not carried away by the flow, an assumption that is compatible with the fact that the typical airspeed of moths and birds largely exceeds the wind velocity. At each time interval Δt, the agent checks if there are odor particles within its olfactory range R_d [see Fig. 1(b)]. If this is the case, then it moves against the direction of the current estimated mean wind at a prescribed speed v_0. When the agent loses contact with the odor cue, it starts the "casting" behavioral program: it moves in a zigzag fashion, always transversally to the current estimated mean wind, with turning times that increase linearly with the time from the last odor detection (see Fig. 1(c) for a sample trajectory and Appendix A for details about the implementation). We denote by $\boldsymbol{v}_i^{\text{priv}}(t)$ the instantaneous velocity of agent i prescribed by this cast-and-surge program. This is uniquely based on private cues, and it would be the actual velocity adopted by the agent when acting in isolation.

B. Response to public cues

To describe the interactions among agents, we have drawn inspiration from flocking and adopted the Vicsek model to describe the tendency of agents to align with their neighbors (see Refs. [21,22], and references therein). We assume that an individual can perceive its peers within a visual range R_a [see Fig. 1(b)] and measure their mean velocity. Within this model, the behavioral response elicited in agent i by its neighbors is

$$\boldsymbol{v}_i^{\text{pub}}(t) = v_0 \sum_{j \in D_i} \boldsymbol{v}_j(t) \left\| \sum_{j \in D_i} \boldsymbol{v}_j(t) \right\|, \tag{2}$$

D_i being the disk of radius R_a centered around the position of the ith individual. As customary, to model imperfection in neighbors' velocity sensing, we added a rotation by a random angle $\boldsymbol{v}_i^{\text{pub}}(t) \leftarrow R(\theta)\boldsymbol{v}_i^{\text{pub}}(t)$, where θ is independently sampled for each agent and at each decision time from a uniform distribution in $[-\eta\pi, \eta\pi]$. The noise strength η ranges from 0 (no noise) to 1 (only noise): in the following we set $\eta = 0.1$. In the absence of external cues and for small enough noise, a group of agents described by this dynamics displays collective flocking and moves coherently in a given direction—unrelated to the source location, however.

C. Combining private and public information

To study the collective olfactory search we merged the two models above, taking the velocity of the ith agent as a linear combination of the two prescriptions arising from private and public cues. The resulting update rule is

$$\begin{aligned} \boldsymbol{v}_i(t) &= (1-\beta)\boldsymbol{v}_i^{\text{priv}}(t) + \beta \boldsymbol{v}_i^{\text{pub}}(t), \\ \boldsymbol{r}_i(t+\Delta t) &= \boldsymbol{r}_i(t) + v_0 \Delta t\, \boldsymbol{v}_i(t)/||\boldsymbol{v}_i(t)||. \end{aligned} \tag{3}$$

The parameter β, which we have dubbed "trust," measures the balance between private and public information. For $\beta = 0$ the agents have no confidence in their peers and behave independently by acting on the basis of the cast-and-surge program only. For $\beta = 1$, agents entirely follow the public cues, discarding any private information, that is, odor detections are completely disregarded.

While it is reasonable to expect that for $\beta = 1$ the unchecked trust in public cues leads to poor olfactory search performances, the nontrivial question here is rather if there is any value at all in public information, i.e., if the optimal β is not equal to zero.

D. Modeling the turbulent environment

To complete the description of our model, we have to specify the flow environment and the ensuing transport of odor particles. In our simulations, the flow environment is given by a time-dependent, incompressible, two-dimensional velocity field, $\boldsymbol{u}(\boldsymbol{x}, t) = \boldsymbol{U} + \boldsymbol{v}(\boldsymbol{x}, t)$, with a constant mean wind $\boldsymbol{U}$ and superimposed statistically stationary, homogeneous, and isotropic fluctuations $\boldsymbol{v}(\boldsymbol{x}, t)$. Odor particles represent patches of odor with concentration above the detection threshold of the agents. The trajectories of odor particles are considered to be tracers and thus follow the dynamics $\dot{\boldsymbol{x}} = \boldsymbol{u}(\boldsymbol{x}, t)$. As for the velocity fluctuations we consider two models: a stochastic flow and a more realistic flow obtained solving the Navier-Stokes equations. Details on the flow, their implementation, and parameters can be found in Appendix B.

III. RESULTS FOR THE STOCHASTIC FLOW

This model flow, characterized by a single length and timescale, is obtained by superimposing a few Fourier modes whose Gaussian amplitudes evolve according to an Ornstein-Uhlenbeck process with specified correlation time. The resulting flow is spatially smooth, exponentially correlated in time, and approximately isotropic (see Appendix B 1 for details).

We studied the performance of collective search as a function of the trust parameter β while keeping the other parameters fixed as detailed in Tables I and II of Appendix C. Initially the agents are waiting in place without any prescribed heading direction until one of the agents detects an odor particle carried by the flow. After this event, agents follow the dynamics (3). Since the search task is a stochastic process, we run many episodes for each value of β to compute the average values of several observables of interest. A given episode is terminated when at least one of the agents is within a distance R_a from the source.

In Fig. 2(a) we show the average time T for the search completion in units of the shortest path time $T_s = L_x/v_0$, which corresponds to a straight trajectory joining the target with the center of mass of the flock at the initial time. There exists an optimal value of the trust parameter $\beta \approx 0.85$ for which agents find the odor source in the quickest way. Remarkably, for this value we obtain $T \simeq 1.03\, T_s$: this means that the agent which arrives first is actually behaving almost as if it had perfect information about the location of the source and were able to move along the shortest path (see Supplemental Material, movie Beta=0.85.mp4 [23]). This result has to be

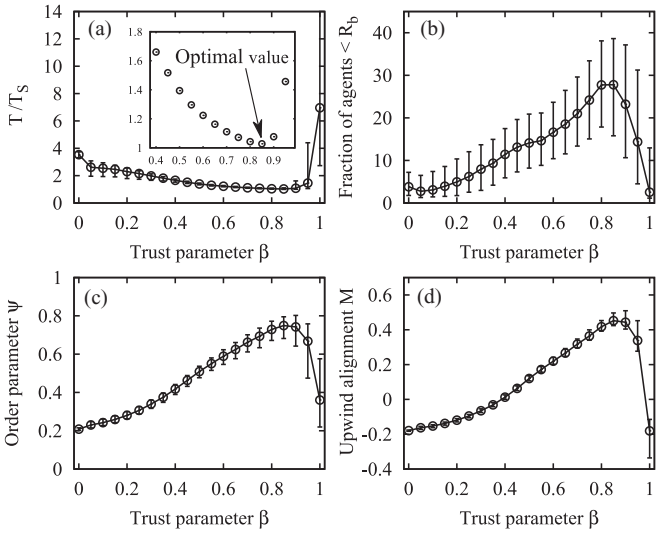

FIG. 2. Collective olfactory search in a stochastic flow. (a) Average search time T for the first agent that reaches the target normalized to the straight-path time, $T_s = L_x/v_0$. The inset enlarges the region close to the minimum. (b) Fraction of agents within a region of size R_b around the source at the first agent arrival time. (c, d) Average order parameter for mutual ψ and wind alignment M. Error bars denote the upper and lower standard deviation with respect to the mean. Statistics is over 10^3 episodes. The parameters are $\lambda = 1$, $N = 100$, $J = 1$, $\eta = 0.1$, $v_0 = 0.5$, $\Delta t = 1$, $L_x = 50$.

contrasted with the singular case of independent agents who act only on the basis of private cues ($\beta = 0$) which display a significantly worse performance (the time to complete the task is more than threefold longer) and move in a zigzagging fashion (see Supplemental Material, movie Beta=0.00.mp4 [23]). It is also important to remark that the average time grows very rapidly as β increases above the optimum. As β approaches unity, agents are dominated by the interactions with their neighbors and pay little attention to odor and wind cues. As a result, they form a flock which moves coherently in an essentially random direction. If, by chance, this direction is aligned against the wind, the task will be completed in a short time. However, in most instances the flock will miss the target and either turn because of the noise η or bounce off the boundaries until, again by sheer chance, some agent will hit the target (see Supplemental Material, movie Beta=0.95.mp4 [23]). This behavior results in a very long average time accompanied by very large fluctuations.

Since we focused on the time of arrival for the first agent reaching the source, it is natural to ask what has happened to the other agents. In Fig. 2(b) we show the average fraction of agents within a distance R_b (the initial size of the group) when the search task is completed, a proxy for the group coherence at arrival. This fraction has a maximum value ≈ 0.3 at about the same value of $\beta \approx 0.85$ that gives the best search time performance. This means that on average about 30% of the group has been moving coherently along the straight path that connects the initial center of mass of the flock to the target.

To quantify the consensus among agents about which direction they have to take, as customary, we introduce the order parameter, $\psi(t) = ||\sum_{i=1}^{N} \boldsymbol{v}_i(t)||/(N v_0)$, with $\psi = 1$ if all the agents move in the same direction and $\psi \sim N^{-1/2} \ll 1$

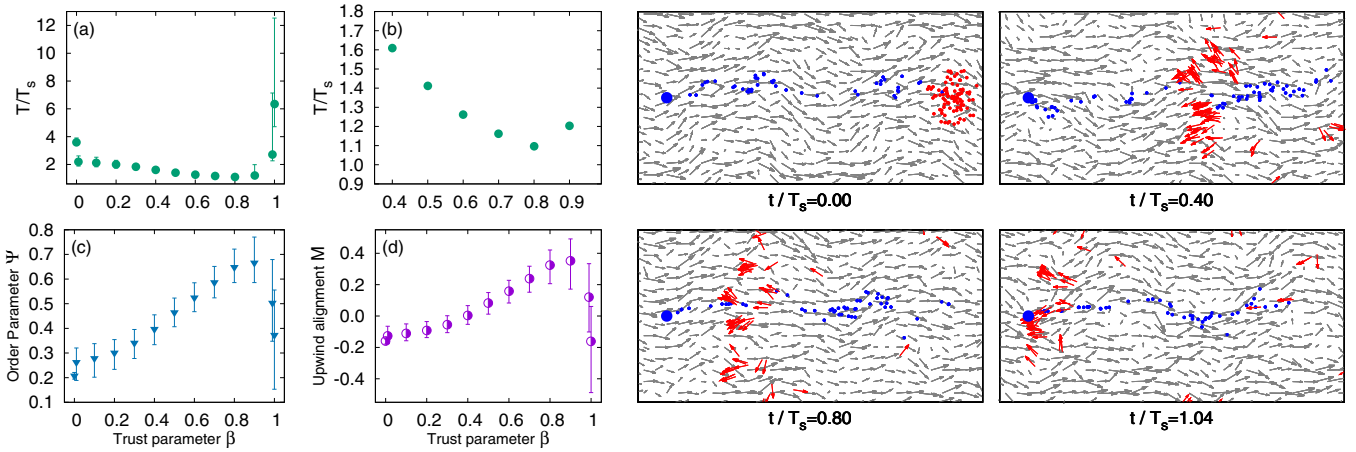

FIG. 3. Collective olfactory search in a turbulent flow. (a) Search time T for the first agent reaching the target normalized by the shortest-path time T_s. (b) Zoom of panel (a) in the region close to the minimum. (c, d) Average order parameter for mutual ψ and wind alignment M. Right panels: four different times t of the search process with the optimal trust parameter $\beta = 0.8$, velocity field (gray arrows), agents (red arrows), odor particles (blue dots), and the source (large blue circle).

if they are randomly oriented. Figure 2(c) shows ψ averaged over all agents and all times. Again we observe a maximum around the values of β where performance is optimal. Another parameter of interest is the upwind alignment of the agents, $M(t) = 1 - \sum_{i=1}^{N} ||\hat{\boldsymbol{U}} + \hat{\mathbf{v}}_i(t)||/N$, where $M = 1$ (-1) means that all the agents are moving upwind (downwind). As shown in Fig. 2(d), the average M has a maximum at around $\beta = 0.85$, which again confirms that a large fraction of the group is heading against the mean wind even if it has access only to a local running time average (the memory time is $\lambda^{-1} = 1 \ll T_s = 100$).

In summary, we found that there is a relatively narrow range of the trust parameter β ($\approx$0.85) for which collective olfactory search is nearly optimal, i.e., the time to reach the target is close to the shortest possible one and takes place with a remarkable group coherence.

Robustness

In order to test the robustness of the above results against variations of the parameters, we have systematically changed (i) the odor particle emission rate J; (ii) the number of agents N; (iii) the range of agent-agent interaction R_a; (iv) agents' speed v_0; and (v) the intensity of flow fluctuations u_{rms}. In order to reduce clutter in the main text, we present the results of this systematic study in a dedicated Appendix, Appendix D, see Figs. 6–10. In summary, we found that the numerical value of the optimal trust parameter is quite robust under a significant variation of parameters [Figs. 6–10 panels (a)]. We also observed an expected loss of coherence for smaller numbers of initial agents or shorter interaction range [Figs. 6–10, panels (b)–(d)]. Interestingly, when the emission rate decreases we observe a deterioration of the performance for agents relying only on private information ($\beta \to 0$), whereas those using public information do not suffer much from the ensuing reduction in the frequency of odor detections. The parameter that has the largest impact on performance is the relative intensity of velocity fluctuations with respect to the mean flow $u_{\rm rms}/U$ [Fig. 10(b)]. When fluctuations increase, the performance significantly degrades, pointing to the key role played by the flow environment in the collective search process.

IV. RESULTS FOR A TURBULENT FLOW

Finally, we considered a more realistic environment where wind fluctuations have a multiscale structure. The flow was obtained by solving the two-dimensional Navier-Stokes equations in the inverse cascade regime [24–26] (see Appendix B 2 for details). As shown in Figs. 3(a)–3(d), the features of multiagent search in a turbulent flow are quite similar to those for the stochastic flow. The average search time [Figs. 3(a) and 3(b)] displays a minimum close to the shortest-path time $T_s = L_x/v_0$ at values of the trust parameter $\beta \approx 0.8$. The right panels of Fig. 3 show four snapshots of the agents at different times during the search process close to optimality. The flock appears to be moving coherently in the upwind direction, and the task is completed in a time $1.04\,T_s$, just a few percent in excess of the nominal minimal time. We remark that the similarity between the results obtained for the stochastic and the turbulent flow (Figs. 2 and 3) is likely due to the specific choice of the sensory input, namely, the agents' decisions are essentially based on single-point and single-time measurements which are not very sensitive to the fine structure of the flow. We predict that agents which can decide on the basis of time-structured inputs (e.g., history of detections) or space-structured (e.g., stereosensing) will show different and distinctive behaviors in single-scale and multiscale flows.

V. CONCLUSIONS AND DISCUSSION

We have shown that there is an optimal way of blending private and public information to obtain nearly perfect performances in the olfactory search task. The first agent that reaches the target completes the task by essentially moving in a straight line to the target. This behavior is striking, since in isolation agents move in a zigzagging fashion [see Fig. 1(c)].

Our results suggest how to build efficient algorithms for distributed search in strongly fluctuating environments.

We remark that the existence of an optimal blending between private and public cues seems to be a feature common to a variety of complex navigational problems. For instance, in the case of pedestrians escaping a smoke-filled room, it has been found that a purely individualistic and purely herding strategies are suboptimal with respect to a mixture of the two [27]. Also, an evolutionary model for migratory populations shows that public information leads to the best migratory routes and the emergence of specialization with the emergence of leaders and followers [28].

The degree of optimal blending between private and public may depend on the specific problem though. Interestingly, in our case, while private information about odor and wind is essential to find the source, its weight in the decision making is numerically rather small, about 15%–20%. This may reflect, at least for the problem of olfactory search, a principle of a "temperate wisdom of the crowds" by which public information must be exploited—but only to a point. We can summarize this as a simple behavioral rule: follow the advice of your neighbors, but once every five to seven times ignore them and act based on your own sensations. It would be interesting in the future to explore the search algorithm that we devised here also in different settings to understand, perhaps even analytically, how the optimal trust parameter depends on the model parameters and the environmental flow properties.

It is important to point out that our algorithm is inherently heuristic. The model heavily draws inspiration from animal behavior, combining features of individual olfactory search in moths and collective navigation in bird flocks. A more principled way of attacking the problem would be to cast it in the framework of multiagent reinforcement learning [29] and seek for approximate optimal strategies under the same set of constraints on the accessible set of actions and on the available private and public information. It would then be very interesting to see if the strategy discovered by the learning algorithms actually resembles the one proposed here, or to other known behaviors displayed by animal groups, or perhaps unveils some yet unknown way of optimizing the integration of public and private cues for collective search.

ACKNOWLEDGMENTS

L.B. acknowledges partial funding from the European Union Programme (FP7/2007-2013) Grant No. 339032. M.D. is grateful for support from the graduate fellowship by the ICTP and University of Trieste. M.D. acknowledges kind hospitality and support from the University of Rome Tor Vergata. M.D. acknowledges fruitful discussions with A. Pezzotta, M. Adorisio, A. Roy, and A. Mazzolini.

M.D. and L.P. contributed equally to this work.

APPENDIX A: DETAILS ON THE ALGORITHMIC IMPLEMENTATION OF THE *CAST-AND-SURGE* SEARCH STRATEGY

The cast-and-surge strategy describes the motion of a single agent elicited by private information in the form of odor encounters and wind estimation. In the following we describe its algorithmic implementation. As discussed in Sec. II, the strategy consists of two components: the estimate of the mean wind velocity $\hat{\boldsymbol{u}}(t)$ and a behavioral response to the presence or absence of an odor within its olfactory range (circle with radius R_d) at a given time. In particular, we assume that the agent can measure the instantaneous local wind at every discrete time δt, which is the integration step used to advance the odor particles. Using such measurements, the agent can construct the estimate of the mean wind velocity $\hat{\boldsymbol{u}}(t)$ by taking an exponentially discounted running average of the perceived flow velocity $\boldsymbol{u}$, i.e., Eq. (1).

Without loss of generality, for the purpose of describing the algorithm we take a simple case where the agent perfectly estimates the mean wind direction at all times (i.e. $\hat{\boldsymbol{u}}(t) = \boldsymbol{U}$). With reference to Eq. (1), this corresponds to the choice $\lambda = 0$ in the memory kernel. Further, we assume that the agent moves every discrete time t separated by the interval $\Delta t \gg \delta t$, here dubbed the decision time. During the time Δt, apart from estimating the mean wind direction every δt, the agent can detect the odor particles within its olfactory range. From a practical perspective, Δt corresponds to the time taken by the agent to make the decision to move by processing the acquired information about the mean wind and the odor detection. Following an extension to continuous space of the cast-and-surge, on-lattice algorithm described by Balkovsky and Shraiman [19], we define the behavioral response of the agent as follows [see Fig. 4(a)]:

Step I: If the agent has detected at least one odor particle in the time interval Δt, it moves upwind by $v_0 \Delta t$ units, v_0 being the speed of the agent. This phase is called "surging." The agent remains in the surging phase as long as it detects odor particles within every Δt time and after taking every step in the surging phase the agent sets $t' = 0$, a number that the agent keeps track of.

Step II: In absence of any odors, the agent moves by $v_0 \Delta t$ units in a direction that forms an angle of $+45^\circ$ with respect to the locally estimated upwind direction.

Step III: The agent updates t' as $t' \leftarrow t' + 2\Delta t$ and then moves in the crosswind direction for time period t' with speed v_0.

Step IV: The agent moves by $v_0 \Delta t$ units in the direction that forms an angle of -45° with respect to the locally estimated upwind direction.

Step V: The agent updates t' as $t' \leftarrow t' + 2\Delta t$ and then moves with speed v_0 in the crosswind direction (opposite to the one taken in step III) for time period t' and resumes further from step II.

Steps II–V describe the "casting" phase, which is terminated as soon as the agent detects an odor particle. Then the agent sets $t' = 0$ and starts the surging phase (step I) from the next decision time.

In Fig. 4(b) we plot a complete sample trajectory of the agent following the cast-and-surge algorithm described above. The ensuing trajectory displays the characteristic zigzag pattern. Two observations are in order. First, the crosswind excursions increase linearly with time. Second, the length traveled in the upwind direction decreases as the inverse square root of time since the last detection. This reflects the fact that the upwind progression is discouraged in the absence of any cues. In the case presented in the main text (for which $\lambda = 1$) the

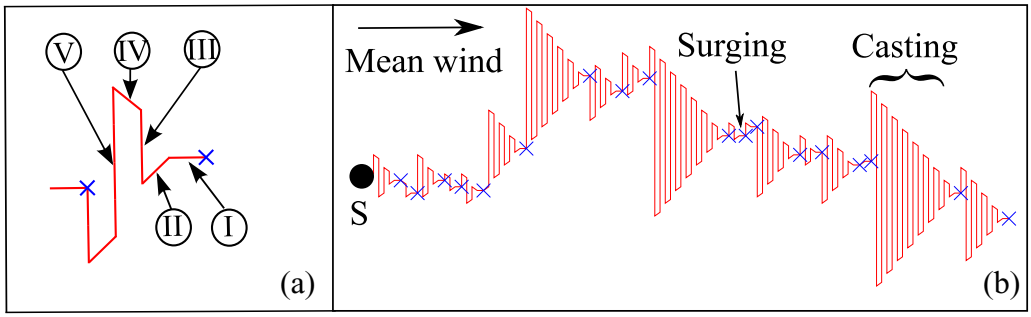

FIG. 4. (a) A short trajectory of an agent navigating according to the cast-and-surge algorithm with $\lambda \to 0$. (b) A complete sample trajectory. The black circle is the location of the source (S), while the blue $\times$'s correspond to the points where it detected an odor particle within its olfactory range.

estimate of the mean wind direction $\hat{\boldsymbol{u}}(t)$ computed by the agent changes with time, as $\lambda > 0$. Thus, in the turbulent environment where the local wind direction fluctuates, the trajectory of the agent deviates from that depicted in Fig. 4(b) as can be see in Fig. 1(c).

APPENDIX B: DESCRIPTION OF THE FLOW ENVIRONMENT

In the following we detail the two models we considered for the fluctuating component of the velocity field.

1. Stochastic flow

As a first simplified setting, we model velocity fluctuations by considering a stochastic flow obtained by superimposing a few Fourier modes, each of them having Gaussian amplitudes, whose real and imaginary part evolve according to independent Ornstein-Uhlenbeck (OU) processes with a specified correlation time τ_f. In this way the resulting flow is spatially smooth and exponentially correlated in time.

Specifically, we consider a flow characterized by a single scale L, obtained by superimposing eight Fourier modes: $\boldsymbol{k} = (k_x, k_y) \in K = K_1 \cup K_2 = \{(k_s, 0), (0, k_s)\} \cup \{(k_s, \pm k_s)\}$, where $k_s = 2\pi/L$. (Notice that we listed only four modes, as the other four are obtained from $\boldsymbol{k} \to -\boldsymbol{k}$, i.e., complex conjugation for maintaining the fields real valued.) The fluctuating velocity is obtained as $\boldsymbol{v}(\boldsymbol{x}, t) = \boldsymbol{\nabla}^{\perp}\psi(\boldsymbol{x}, t)$ with $\boldsymbol{\nabla}^{\perp} = (-\partial_y, \partial_x)$, and the stream function ψ is computed at each odor particle position by means of the following formula:

$$\psi(\boldsymbol{x}, t) = \sum_{\boldsymbol{k}\in K} (A(\boldsymbol{k}, t)e^{i\boldsymbol{k}\cdot\boldsymbol{x}} + \text{c.c.}), \tag{B1}$$

where c.c. stands for the *complex conjugate*. The amplitudes of the Fourier modes $A(\boldsymbol{k}, t)$ are Gaussian random complex variables evolving with the following OU process:

$$\partial_t A_\gamma(\boldsymbol{k}, t) = -\frac{1}{\tau_f} A_\gamma(\boldsymbol{k}, t) + \left(\frac{2\sigma^2(\boldsymbol{k})}{\tau_f}\right)^{\frac{1}{2}} \eta_\gamma(\boldsymbol{k}, t), \tag{B2}$$

where γ labels the real and imaginary part, and $\eta_\gamma(\boldsymbol{k}, t)$ are zero mean Gaussian variables with correlation $\langle \eta_\gamma(\boldsymbol{k}, t)\eta_{\gamma'}\boldsymbol{k}', t')\rangle = \delta_{\gamma,\gamma'}\delta_{\boldsymbol{k},\boldsymbol{k}'}\delta(t - t')$ and so that

$$\langle A_\gamma(\boldsymbol{k}, t)A_{\gamma'}(\boldsymbol{k}', t')\rangle = \sigma^2(\boldsymbol{k})\delta_{\gamma,\gamma'}\delta_{\boldsymbol{k},\boldsymbol{k}'} \exp(-|t - t'|/\tau_f).$$

The standard deviations $\sigma(k)$ have been chosen to have an approximately isotropic velocity field with full control on the fluctuation intensity $u_{\text{rms}} = \langle (v_x^2 + v_y^2)/2\rangle^{1/2}$. In particular, we take $\sigma(k) = cu_{\text{rms}}/(\sqrt{3}k_s)$, with $c = 1$ for $\boldsymbol{k} \in K_1$ and $c = 1/2$ for $\boldsymbol{k} \in K_2$ so that $\langle v_x^2\rangle = \langle v_y^2\rangle = u_{\text{rms}}^2$.

Similar flows have been used for studying, e.g., the statistical dynamics of inertial particles [30,31]. In our simulations, the constant mean wind is fixed to $U = 1$ and the fluctuation intensity to $u_{\text{rms}} = 0.42U$. As for the fluctuating component, it has one single characteristic scale set to $L = 10$, and a correlation time of the amplitudes of the Fourier modes equal to $\tau_f = 5$.

Tests about the search conducted by one single agent have been done considering different values of the flow parameters, also introducing more than one scale. Such tests have shown the same qualitative behaviors reported here, provided that u_{rms} remains smaller than U.

2. Turbulent flow

As a more realistic case we consider velocity fluctuations obtained from a direct numerical simulation (DNS) of the two-dimensional Navier-Stokes equations (NSEs) in the inverse cascade regime [24–26]. The NSE written for the vorticity field, $\omega = \boldsymbol{\nabla} \times \boldsymbol{u}$, reads

$$\partial_t\omega + \boldsymbol{v} \cdot \boldsymbol{\nabla}\omega = \nu\Delta\omega - \alpha\omega + f, \tag{B3}$$

where $\boldsymbol{v} = \nabla^{\perp}\psi(\boldsymbol{x}, t)$, where the stream function ψ is obtained by inverting $\omega = -\Delta\psi$. The DNS of Eq. (B3) was carried out using a standard 2/3 dealiased pseudospectral solver over a biperiodic $2\pi \times 2\pi$ box with second-order Runge-Kutta time stepping. Energy and enstrophy are injected at rates ϵ and ζ, respectively, by the forcing term f which is a zero mean, Gaussian field with correlation $\langle f(\boldsymbol{x}, t)f(\boldsymbol{0}, t')\rangle = \delta(t - t')F(r/\ell_f)$ acting at small scales, $\ell_f \ll 2\pi$, with $F(x) = F_0\ell_f^2 \exp(-x^2/2)$. With this forcing, an inverse energy cascade sets in at scales $r \gg \ell_f$. In order to establish a statistically steady state, the Ekman friction term, $-\alpha\omega$, extracts energy at large scales, $L_\alpha \approx \epsilon^{1/2}\alpha^{-3/2}$, while the viscous term removes enstrophy at small scales. As a result, we have a velocity field which is nonsmooth in the inertial range of scales, $\ell_f \ll r \ll L_\alpha$, and smooth below ℓ_f. In Fig. 5 we show the mean energy spectrum $E(k)$ displaying the Kolmogorov $k^{-5/3}$ scaling behavior, which means that

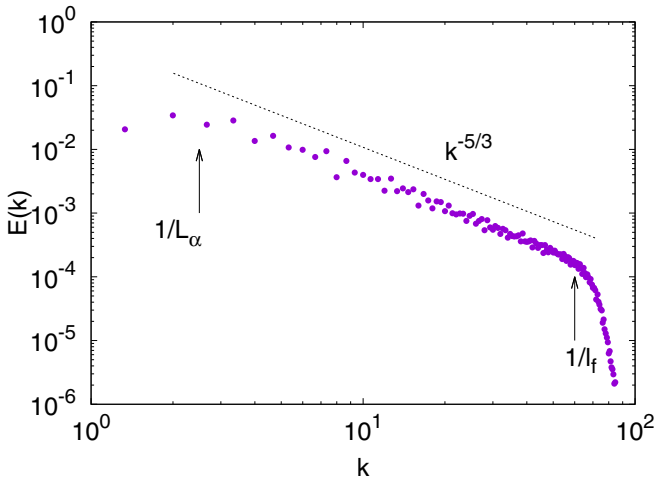

FIG. 5. Energy spectrum obtained by direct numerical simulations of Eq. (B3) with 256^2 grid points. Hyperviscous dissipation of order 8 has been used with viscosity $\nu_8 = 1.3 \times 10^{-29}$, Ekman friction coefficient $\alpha = 0.02$, and time step $dt = 10^{-3}$. The large scale of the flow is about half of the simulation box.

in the inertial range velocity differences over a scale r are approximately Hölder continuous with exponent $1/3$.

Owing to the necessity of storing the entire history of the full velocity field (see below for details), we used a relatively small resolution of 256^2 grid points. Thus, to reduce as much as possible the enstrophy cascade range, we used a hyperviscous term of order 8, which removes enstrophy very close to the injection scale; this is a customary procedure when interested in simulating the inverse cascade in low-resolution DNS (see Refs. [24,25]).

In order to evolve the odor particles and perform statistics over many episodes of the collective search, we stored the whole evolution of the velocity field for about ten large-scale timescales, $T_{L_\alpha} \approx 5$. The velocity field history is then cycled in time, so the flow is effectively periodic in time with a period of about $10T_{L_\alpha}$. For each episode we place the source in a different position within the simulation box and define the mean wind direction to be either along the horizontal or vertical direction (this is done to average over different flow regions). We let the source emit the particles at exponentially distributed times with average $\tau = 5$, which corresponds to the timescale associated to the forcing scale, and advect them in the full plane (making use of the spatial periodicity), with a velocity obtained by interpolating the velocity field at the particle position and superimposing the mean wind U. We wait until the statistics of the odor particles becomes stationary in the region of interest and then let the searching agents look for the source.

Then the agents are initially placed at distance L_x downwind from the source [see Fig. 1(a)] and wait for the first detection to start the search. The episode ends when one of the agents reaches the source as described in main text.

APPENDIX C: TABLE OF PARAMETERS

Here we summarize the main parameters defining the cast-and-surge algorithm (Table I) and the flow (Table II) used in the main text. It is worth pointing out that the stochastic and turbulent flow have different characteristic length and timescales. As a consequence, the decision time and the detection radius in the two flows are different (see Table II). Thus, in order to study the olfactory search in comparable regimes, we rescaled all the other quantities accordingly, maintaining at the same time identical ratios among them.

TABLE I. Values of the parameters of the cast-and-surge algorithm. The dimensional quantities are written in terms of the radius of detection of one agent R_d and its decision time Δt, which are detailed in Table II.

Description	Symbol	Numerical Value
Initial distance between source and agents' center of mass	L_x	$250R_d$
Simulation box size factor	b	2.5
Number of agents	N	100
Emission rate of odors from the source	J	1.0 particle/Δt
Initial cluster size of agents	R_b	$25R_d$
Range of agent-agent interaction	R_a	$5.0R_d$
Speed of the agents	v_0	$2.5R_d/\Delta t$
Strength of the noise	η	0.1
Inverse of the memory time	λ	$1.0/\Delta t$
Mean wind intensity	U	1.0

APPENDIX D: DEPENDENCE ON VARIOUS PARAMETERS

In the main text we presented results on the multiagent search with parameters fixed as in Tables I and II. In this Appendix, we carry out a systematic study at varying the model parameters, within the stochastic flow environment detailed in Appendix B 1. We studied the performance of the multiagent system by varying one parameter at a time, in particular, we considered variations of the emission rate of odor J, the number of agents N, range of agent-agent interaction R_a, speed of the agents v_0, and the flow fluctuation intensity $u_{\rm rms}$. The main findings are summarized in Sec. III, while the next sections detail the results at changing the above listed five parameters. In order to make the comparison as clear as possible, in all subsequent figures data plotted in black colors refer to the case shown in the main text.

1. Dependence on the emission rate of odor J

We carried out simulations with various emission rates J, i.e., the rate at which odor particles are released by the

TABLE II. Values of the agent radius of detection R_d and its decision time Δt in each flow configuration and descriptions of the flow characteristics.

Description	Stochastic flow	Turbulent flow
Decision time Δt	1.0	0.2
Olfactory range of the agent R_d	0.2	0.04
Fluctuation intensity $u_{\rm rms}$	0.42U	0.42U
Characteristic length	10.0	2.0
Characteristic time	5.0	5.0

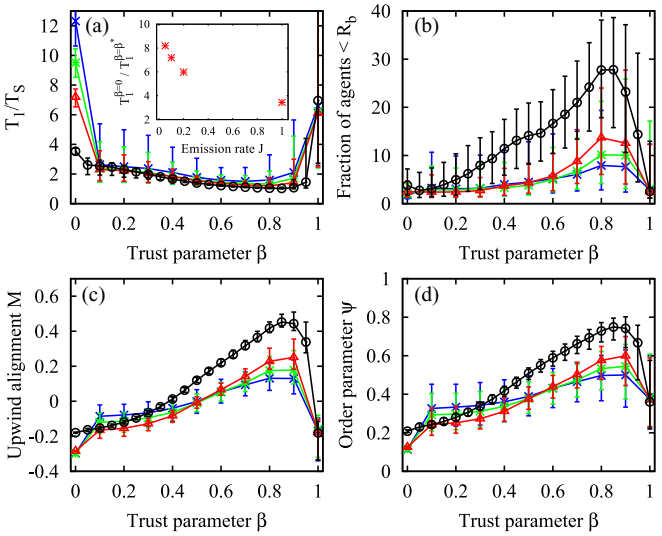

FIG. 6. Collective olfactory search with various emission rates J, other parameters as in Tables I and II: (black circles) $J = 1$, (red triangles) $J = 0.2$, (green stars) $J = 0.1$, (blue crosses symbols) $J = 0.05$. (a) Average search time T_1 normalized to $T_s = L_x/v_0$. The inset shows the ratio of average time taken by agents with only private information ($T_1^{\beta=0}$) to the average time taken by agents with optimal integration of private and public information ($T_1^{\beta=\beta^*}$). (b) Fraction of agents within a region of size R_b around the source at the arrival time of the first agent reaching the target. (c) Average alignment against the mean wind M. (d) Average order parameter ψ. For all data, the error bars denote the upper and lower standard deviation with respect to the mean value. Statistics is over 10^3 episodes.

source. As J is decreased, the odor signal detected by the agents becomes more intermittent and, consequently, locating the source more challenging.

In Fig. 6 we show the quantities of interest for different values of J as a function of the trust parameter β. Figure 6(a) shows that at decreasing J, the time to locate the source

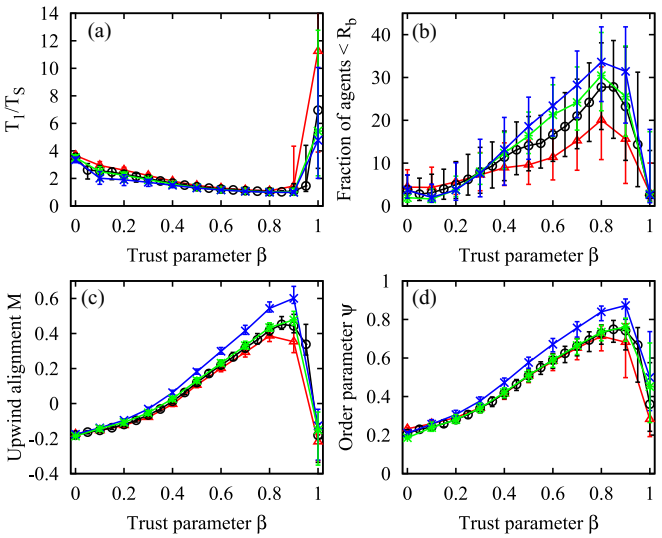

FIG. 7. Collective olfactory search for different values of the number of agents N, other parameters as in Tables I and II: (red triangles) $N = 50$, (black circles) $N = 100$, (green stars) $N = 150$, (blue crosses symbols) $N = 200$. Panel description as in Fig. 6.

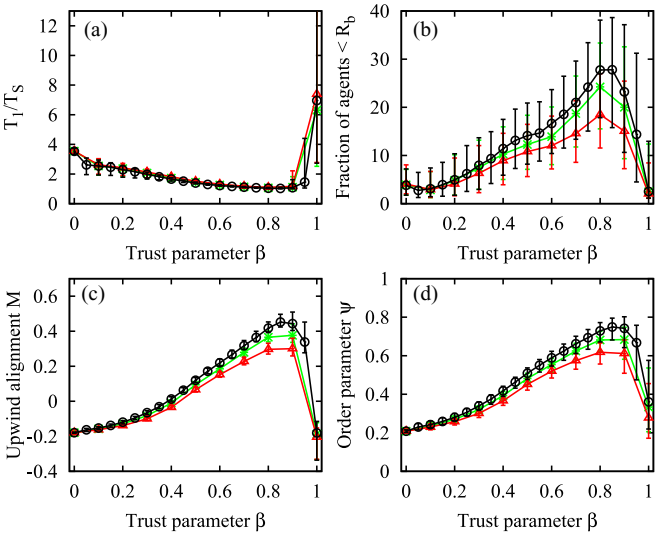

FIG. 8. Collective olfactory search for different values of the agent-agent interaction range, other parameters as in Tables I and II: (black circles) $R_a = 1.0$, (green stars) $R_a = 0.85$, (red triangles) $R_a = 0.70$. Panel description as in Fig. 6.

by agents using only private information ($\beta = 0$) increases. However, with optimal combination of public and private information (i.e., optimal trust parameter $\beta^* \approx 0.8-0.85$) the agents could locate the odor source more or less in the same minimum amount of time for various emission rates J. This means that agents optimally combining public and private information can buffer a sparse odor signal better than agents relying only on private information. The relative performance characterized as the ratio of average time taken to locate the source by agents with $\beta = 0$ and agents with optimal value of $\beta = \beta^*$ is shown in the inset of Fig. 6(a). Figure 6(b) shows that the fraction of agents within the distance R_b from the

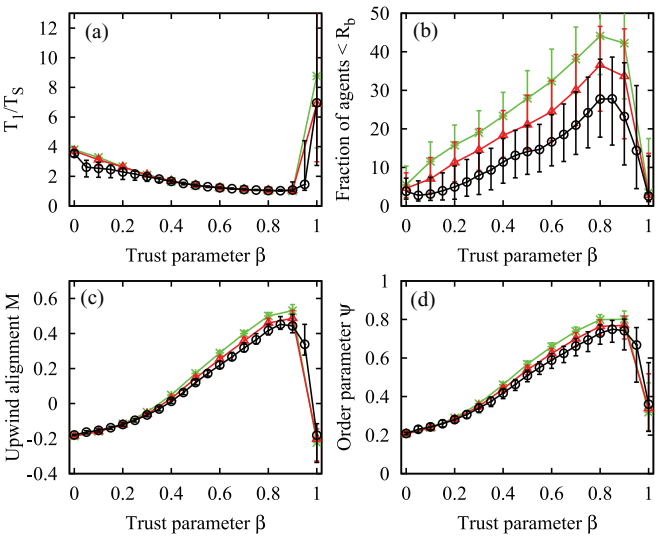

FIG. 9. Collective olfactory search with different agents' speeds, v_0, other parameters as in Tables I and II: (black circles) $v_0 = 0.50$, (red triangles) $v_0 = 0.25$, (green stars) $v_0 = 0.125$. Panel description as in Fig. 6.

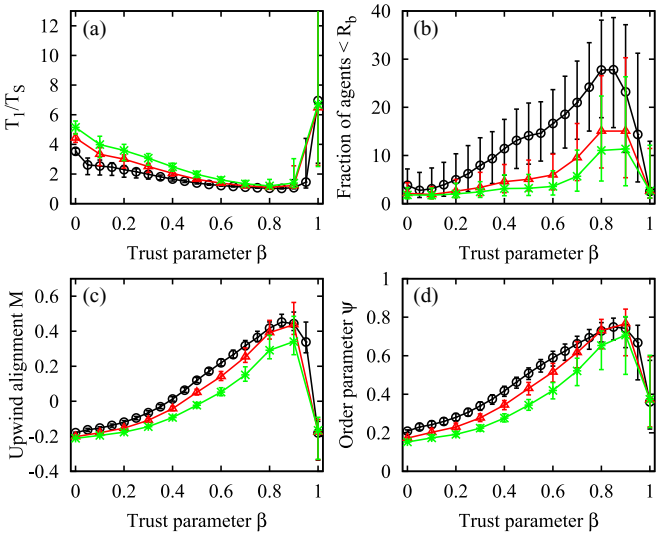

FIG. 10. Collective olfactory search with various intensities of fluctuations in the model flow, other parameters as in Tables I and II: (black circles) $u_{\rm rms} = 0.4242U$, (red triangles) $u_{\rm rms} = 0.6363U$, (green stars) $u_{\rm rms} = 0.8484U$. Panel description as in Fig. 6.

source at the moment the source is located has a maximum for $\beta = \beta^*$. Note that the fraction of agents near the source with the optimal trust parameter β^* depends strongly on the emission rate of odor. In Figs. 6(c)–6(d) we show the average upwind alignment and polar order parameter, defined in the main text. Both quantities peak for $\beta = \beta^*$, consistently with the general results presented in the main text: the agents with optimal value of trust parameter have the highest direction consensus, and most of the time they move against the mean wind M.

2. Dependence on the number of agents

We now vary the number of agents while keeping constant their initial density by changing the initial cluster size R_b of the agents appropriately. In Fig. 7 we show that results do not depend in any important way on N except for the obvious case when we measure the number of agents that are able to coherently reach the target [Fig. 7(b)], which decreases by decreasing the total number of agents.

3. Dependence on range of agent-agent interaction

In this section we consider the effect of the variation of the agent-agent interaction range, R_a. Figure 8 summarizes the results. Again, we observe that upon varying R_a most of the results are robust, except for the number of agents able to reach the target simultaneously, which decreases by decreasing the interaction range as shown in panel (b).

4. Dependence on the speed of the agents

Here we study the impact of different agent speeds v_0 on performance. As shown in Fig. 9, also in this case we observe robustness for the optimal trust parameter value, panel (a). Clearly the absolute time to reach the target for the first agent increases at decreasing v_0 (this cannot be appreciated in panel a due to normalization with the time taken for a straight path). However, we observe an increase in the coherence of the group [Fig. 9(b)] at decreasing the agent speed v_0, which is reasonably explained by the fact that by going slower they can remain closer to each other. This suggests that if arriving at the target in a large number more or less at the same time is important, using slower agents may be a good strategy.

5. Dependence on flow with various intensities of fluctuation

In this last section we describe the dependency on the most sensitive parameter, which is the intensity of the flow fluctuations defined in terms of $u_{\rm rms}/U$, where U is the amplitude of the mean flow. As expected, by increasing the environment *noise* everything becomes more complicated; the number of agents at the target is sensibly smaller [Fig. 10(b)], as well as the alignment and the order parameter [Figs. 10(c) and 10(d)]. The value of the optimal trust β^* is nevertheless strongly robust, as shown in Fig. 10(a).

[1] C. David, J. Kennedy, and A. Ludlow, Nature (London) **303**, 804 (1983).
[2] J. Kennedy, Physiol. Entomol. **8**, 109 (1983).
[3] J. Elkinton, C. Schal, T. Ono, and R. Cardé, Physiol. Entomol. **12**, 399 (1987).
[4] T. Dekker, M. Geier, and R. Cardé, J. Exp. Biol. **208**, 2963 (2005).
[5] J. Murlis and C. Jones, Physiol. Entomol. **6**, 71 (1981).
[6] A. Celani, E. Villermaux, and M. Vergassola, Phys. Rev. X **4**, 041015 (2014).
[7] A. Lilienthal, D. Reimann, and A. Zell, Autonome Mobile Systeme **2003**, 150 (2003).
[8] G. Ferri, E. Caselli, V. Mattoli, A. Mondini, B. Mazzolai, and P. Dario, in *Proceedings of the First IEEE/RAS-EMBS International Conference on Biomedical Robotics and Biomechatronics, Pisa, Italy* (IEEE, New York, 2006), p. 573.
[9] H. Ishida, H. Tanaka, H. Taniguchi, and T. Moriizumi, Auton. Robots **20**, 231 (2006).
[10] F. Grasso and J. Atema, J. Environ. Fluid Mech. **2**, 95 (2002).
[11] T. Lochmatter and A. Martinoli, Exper. Robot. **54**, 473 (2009).
[12] E. Bonabeau, G. Theraulaz, and M. Dorigo, *Swarm Intelligence: From Natural to Artificial Systems* (Oxford University Press, New York, 1999).
[13] T. Pitcher, A. Magurran, and I. Winfield, Behav. Ecol. Sociobiol. **10**, 149 (1982).
[14] C. Torney, Z. Neufeld, and I. Couzin, Proc. Nat. Acad. Sci. USA **106**, 22055 (2009).
[15] C. Ioannou, M. Singh, and I. Couzin, Am. Nat. **186**, 284 (2015).
[16] A. Berdahl, C. Torney, C. Ioannou, J. Faria, and I. Couzin, Science **339**, 574 (2013).
[17] N. Miller, S. Garnier, A. Hartnett, and I. Couzin, Proc. Nat. Acad. Sci. USA **110**, 5263 (2013).

[18] J.-B. Masson, M. Bailly, and M. Vergassola, J. Phys. A **42**, 434009 (2009).
[19] E. Balkovsky and B. Shraiman, Proc. Nat. Acad. Sci. USA **99**, 12589 (2002).
[20] T. Baker, M. Willis, and P. Phelan, Physiol. Entomol. **9**, 365 (1984).
[21] T. Vicsek and A. Zafeiris, Phys. Rep. **517**, 71 (2012).
[22] F. Ginelli, Eur. Phys. J. Spec. Top. **225**, 2099 (2016).
[23] See Supplemental Material at http://link.aps.org/supplemental/10.1103/PhysRevE.102.012402 for the movies discussed in the text.
[24] G. Boffetta, A. Celani, and M. Vergassola, Phys. Rev. E **61**, R29 (2000).
[25] G. Boffetta and R. E. Ecke, Annu. Rev. Fluid Mech. **44**, 427 (2012).
[26] A. Alexakis and L. Biferale, Phys. Rep. **767–769**, 1 (2018).
[27] D. Helbing, I. Farkas, and T. Vicsek, Nature (London) **407**, 487 (2000).
[28] C. J. Torney, S. A. Levin, and I. D. Couzin, Proc. Nat. Acad. Sci. USA **107**, 20394 (2010).
[29] L. Buşoniu, R. Babŭska, and B. Schutter, IEEE Trans. Syst., Man, Cybern. C **38**, 156 (2008).
[30] J. Bec, Phys. Fluids. **15**, L81 (2003).
[31] J. Bec, A. Celani, M. Cencini, and S. Musacchio, Phys. Fluids. **17**, 073301 (2005).