

Optimal tracking strategies in a turbulent flow

Chiara Calascibetta^{1,2} [✉], Luca Biferale^{1,2} , Francesco Borra³, Antonio Celani⁴  & Massimo Cencini^{2,5} 

Pursuing a drifting target in a turbulent flow is an extremely difficult task whenever the searcher has limited propulsion and maneuvering capabilities. Even in the case when the relative distance between pursuer and target stays below the turbulent dissipative scale, the chaotic nature of the trajectory of the target represents a formidable challenge. Here, we show how to successfully apply optimal control theory to find navigation strategies that overcome chaotic dispersion and allow the searcher to reach the target in a minimal time. We contrast the results of optimal control – which requires perfect observability and full knowledge of the dynamics of the environment – with heuristic algorithms that are reactive – relying on local, instantaneous information about the flow. While the latter display worse performances, optimally controlled pursuers can track the target for times much longer than the typical inverse Lyapunov exponent and are considerably more robust.

¹Department of Physics, University of Rome “Tor Vergata”, Via della Ricerca Scientifica 1, 00133 Rome, Italy. ²INFN, Sezione di Roma “Tor Vergata”, Rome, Italy. ³Laboratory of Physics of the École Normale Supérieure, 24 Rue Lhomond, 75005 Paris, France. ⁴Quantitative Life Sciences, The Abdus Salam International Centre for Theoretical Physics, ICTP, 34151 Trieste, Italy. ⁵Istituto dei Sistemi Complessi, CNR, Via dei Taurini 19, 00185 Rome, Italy. [✉]email: calascibetta@roma2.infn.it

Finding optimal navigation strategies in a complex fluid environment is a notoriously difficult problem with applications ranging from environmental monitoring^{1–4} to micro-medicine^{5–7}. A well-explored set of navigation problems is point-to-point path-planning optimization of flying vehicles such as airplanes or drones, with the aim of minimizing some functioning cost that may comprise fuel consumption and time of arrival^{8–10}. These vehicles move in a complex chaotic environment but can have almost full control on their trajectory as their speed is typically larger than the fluid velocity. Recently, point-to-point path-planning optimization has been the focus of intense research also for microswimmers and active particles. These slow, microscopic objects tend to be carried away by the flow and need to appropriately exploit it in order to reach their destination^{11–16}. In Ref. 17 it has been studied how the optimal path is influenced by hydrodynamics interactions, while in^{18,19} the importance of harnessing vortical structures is highlighted. Ref. 20,21 considers a predation problem at low Reynolds number, while in^{22,23} the optimal navigation of colloidal robots has been investigated. In^{24,25} the authors have analyzed geometric approaches for optimization, while in^{26,27} multi-objective reinforcement learning methods have been applied.

Here, we consider a more difficult navigation task, where the target is not fixed in space but is chaotically advected by the turbulent flow. The challenge of tracking a Lagrangian target is increased by the limited speed and manoeuvrability of the pursuer. This problem is relevant to many applications, such as keeping in a pattern formation a swarm of oceanic drifters and floaters^{9,28–31}; interpreting strategies to catch non-swimming preys by micro-swimmers in turbulent environment; developing autonomous self-propelling protocols for mini-robots navigating in complex bio-flows or for gliders in the atmosphere or oceans^{3,32–39}. We will consider the case in which the searcher and target stay at distances smaller than the Kolmogorov length and thus experience a smooth chaotic flow at all times. As depicted in Fig. 1a, due to its limited speed the agent must know how to *surf* local eddies in ingenious ways, by taking advantage of strong fluctuations, sometimes generated by vortical structures (Fig. 1 panels b–c), and exploiting the long time correlations typical of turbulence (Fig. 1d). The final objective of the optimal pursuer is to catch the chaotic moving target in the shortest possible time or, if all else fails, to be as close as possible to it at the end of the allotted time for the pursuit.

The main obstacle in finding the optimal control is that it depends on the entire spatio-temporal history of the chaotic system evolution. Analytical solutions can be found only in simple linear or time invariant flows, limited to the point-to-point navigation task^{17,40–43}. An interesting way to explore the problem in a simplified way, would be to remap it into a biased random walk. Despite the absence of non-Gaussian and intermittent statistics, this could provide valuable theoretical insights by allowing to recover analytically the probability of reaching the target and enabling a priori performance evaluation. In this paper, we show how to apply Optimal Control Theory to discover the best controls for the turbulent tracking problem and compare them with heuristic reactive strategies. In particular, since the problem is formally akin to a pursuit game we use for comparison heuristics like the (i) pure pursuit⁴⁴, i.e. always swimming along the line-of-sight of the target, and (ii) two other heuristic controls obtained by solving simplified optimality problems over a short time horizon with the help of some assumptions on the dynamics. We show that such heuristic strategies, which certainly have the advantage of not requiring a large amount of information on the system, are markedly sub-optimal. How to develop effective heuristic strategies with limited information and computation but whose

performances are closer to optimality remains an outstanding problem.

Results

Tracking Lagrangian targets. The pursuer—hereafter also called the agent—and the target are immersed in a turbulent velocity field $\mathbf{u}_t(\mathbf{x})$ obtained by direct numerical integration of the Navier-Stokes equation, sustained by an isotropic and homogeneous forcing, with Reynolds number at the Taylor scale $Re_\lambda \simeq 130$ ^{45,46} (see “Methods” for details). The flow is statistically stationary and we will consider a single realization of the time-dependent velocity field from the initial time 0 until the final time, t_f . Each episode corresponds to a different initial position of the target and thus a different history of the velocity field along the trajectory of the tracer. Episodes start with the target, $\mathbf{X}_0$, and the agent, $\mathbf{X}_t^{(a)}$, placed at a distance of the order of the Kolmogorov scale, $|\mathbf{X}_0^{(a)} - \mathbf{X}_0| = |\mathbf{R}_0| = R_0 \simeq \eta$. The flow transports both the target, which moves as a passive tracer, and the agent that, however, can exert some control by swimming, with velocity $\mathbf{U}_t$, with respect to the medium. While it can freely point in any direction $\hat{\mathbf{n}}_t$, the swimming velocity is constrained to have a fixed speed, V_s . We chose this speed to be smaller than the Kolmogorov velocity, u_η . Specifically, the results shown below are for $V_s \simeq 0.13u_\eta$.

Target and agent therefore move according to the dynamics

$$\begin{cases} \dot{\mathbf{X}}_t = \mathbf{u}_t(\mathbf{X}_t) \\ \dot{\mathbf{X}}_t^{(a)} = \mathbf{u}_t(\mathbf{X}_t^{(a)}) + \mathbf{U}_t \\ \mathbf{U}_t = V_s \hat{\mathbf{n}}_t, \end{cases} \quad (1)$$

where we have assumed that the time scale associated with the agent to align with the flow is small compared to τ_η . In other words, we consider the reorienting torque strong enough such that the agent reorients instantaneously to the steering protocol. The episode lasts up to a maximum time $t_f \simeq 50\tau_\eta$, i.e. much larger than the inverse of the Lagrangian Lyapunov exponent of the tracer trajectory $\lambda^{-1} \simeq 7.5\tau_\eta$ ⁴⁷, where τ_η is the Kolmogorov time. Within this time horizon the agent has to find the optimal choice of the steering protocol $\hat{\mathbf{n}}_t$ that allows to capture the target. The capture event is defined as the moment when the relative distance becomes smaller than the capture distance $R_c = 10^{-2}R_0$, which is much smaller than the initial separation. From that time on, the agent sticks to the target. The speed is set as $V_s \approx \lambda R_0$ so that the agent is at the limit of controllability, i.e. the control is smaller than the typical background velocity between the agent and target.

As long as the agent and target remain within distances where the flow is differentiable (i.e. $R_t \leq O(10\eta)$, where R_t is the agents distance at the time t and 10η is the order of the correlation scale of the velocity gradients^{48,49}) the fluid velocity experienced by the searcher is approximately $\mathbf{u}_t(\mathbf{X}_t^{(a)}) \simeq \mathbf{u}_t(\mathbf{X}_t) + \nabla \mathbf{u}_t \mathbf{R}_t$, where $\nabla \mathbf{u}_t$ is the flow gradient evaluated at the target position. As a result, the optimization problem depends only on the separation vector $\mathbf{R}_t$ and on the entire history of the velocity gradients along the trajectory of the searching agent. We will be enforcing this approximation, which allows us to store Lagrangian trajectories and their accompanying velocity gradient to create a database that can be used to compute the optimal control (see “Methods”).

Heuristic control strategies. As a term of comparison, we will consider several heuristic controls that neither require knowledge of the future evolution of the trajectory (required for the computation of the optimal control) nor the memory of its past. These are called reactive strategies as they are purely based on instantaneous information about the flow and the pair separation.




Fig. 1 Graphical illustration of the system setup. **a** A pursuing agent (in red color) with limited maneuverability aims to stay as close as possible or, possibly, to capture a Lagrangian target (in yellow color) chaotically advected by a turbulent flow. The background is given by a rendering of the turbulent vorticity intensity at a fixed time during the episode. **b** Time evolution of the typical dimensionless transverse ($\partial u_z / \partial y$ in black color), and longitudinal ($\partial u_z / \partial z$ in red color) velocity gradients during the duration of a catching episode. These quantities are related to vorticity and shear stress, respectively. Since the flow is isotropic the two components above are sufficient to describe the statistical properties of the flow. The normalization factor is given by the Kolmogorov time-scale of the flow, $\tau_\eta = \sqrt{\nu/\epsilon}$, where ν is the fluid viscosity and ϵ the mean energy dissipation. **c** Probability distribution function (PDF) of the transverse and longitudinal gradients measured over 2×10^5 trajectories of length $T \simeq 150\tau_\eta$. **d** Time correlation functions for the transverse and longitudinal velocity gradients.

1. *Pure Pursuit (PP)*. The agent only knows in which direction the target currently is and swims towards it⁴⁴:

$$\hat{\mathbf{n}}_t^{\text{PP}} = -\hat{\mathbf{R}}_t.$$

2. *Surfing Control (SC)*. The agent has knowledge of the instantaneous velocity gradient and direction of the target^{50,51}:

$$\hat{\mathbf{n}}_t^{\text{SC}} \propto -[\exp(\tau_s \nabla \mathbf{u}_t)]^\top \hat{\mathbf{R}}_t,$$

where τ_s is a parameter to be chosen empirically. This control maximizes the displacement of the pursuer towards the target in a time interval τ_s , assuming that the gradient and the direction of the target $\hat{\mathbf{R}}_t$ are constant over that time (see “Methods”).

3. *Perturbative Optimal Control (PO)*. The agent has the same information as in (SC) but the direction in which it swims is

$$\hat{\mathbf{n}}_t^{\text{PO}} \propto -[\exp(\tau_p \nabla \mathbf{u}_t)]^\top \exp(\tau_p \nabla \mathbf{u}_t) \hat{\mathbf{R}}_t.$$

As discussed in “Methods”, we obtained this control strategy by solving perturbatively the optimal control problem which minimizes the distance from the target after a time τ_p , assuming constant gradients (but dynamically evolving separations). The time τ_p is a free parameter to be chosen empirically.

Optimal control. How can one define the performance of a tracking agent? Ideally, the searcher should be able to capture the moving target in the shortest possible time, and in any event

before a certain time horizon. However, this is not always possible because of the strong dispersion induced by the underlying turbulent flow. In these situations one could instead settle for the less ambitious objective of reaching the shortest possible separation at the end of the allotted time for the search. These requests can be summarized by the following mathematical expression for the cost function

$$J = R_{t_f}^2 + c\lambda R_c^2 \int_0^{t_f} dt \Theta(R_t - R_c). \quad (2)$$

The first term of the cost function J is the distance between agent and target at the final time while the second one, Θ being the Heaviside function, is the time needed to reach the capture distance, R_c . If the episode ends with a capture for $t < t_f$ the first term equals R_c^2 and does not depend on the control so that the task is to minimize the time for capture. Conversely, if the capture does not occur before t_f the second term is always equal to $c\lambda R_c^2 t_f$ regardless of the trajectory and the task reduces to minimizing the final separation R_{t_f} . The constant c controls the trade-off between the two contributions and it is chosen to give a strong enough weight to the capture events.

The theory of Optimal Control (OC) provides the tools to compute the best control – i.e. the one with minimal cost – by casting the optimization in the form of Euler-Lagrange equations that have to be solved both forward and backward in time, requiring full knowledge of the dynamics of the system (see “Methods”) ^{40,52,53}. If a solution to the extremality conditions exists and is stable, it provides at each time the optimal direction




Fig. 2 Capture time statistics. **a** PDF of the capture time for the Optimal Control strategy (red open circles) and for the heuristic strategies, Pure Pursuit (black full rhombus), Surfing Control (blue full triangles) and Perturbative Optimal control (green full squares). The PDFs are evaluated along the episodes where all strategies capture. The vertical dashed line represents the time horizon, t_t . In the inset we show the same PDF in log-scale. **b** Box-plot of the catching time PDF. The box shaded area reports the range 25th–75th percentile, the solid and dashed lines are the median and the mean respectively. Each whisker contains the remaining 25% of the data, while outliers points are identified for a large value of the capture time. **c** PDF of the capture time for the reactive heuristic strategies, normalized with the corresponding capture time for the OC, $t_c/t_c^{(OC)}$ constrained to the episodes where all the heuristic strategies capture within t_t .

$\hat{n}_t^{OC}$ which implicitly depends on the whole history—past, present and future – of gradients.

Capture time statistics. We sampled 2×10^5 Lagrangian target trajectories with the agent starting at a random position distant $R_0 \simeq \eta$ from the target. The OC algorithm converges in 81% of the cases. Among the optimally controlled paths, the target is successfully captured in 81% of the cases. With further fine-tuning of the hyper-parameters involved in the iterative process, one might still improve the convergence of the algorithm up to probably 100%. Over the same converged set, the percentages of success of the three heuristic strategies are similar and close to 70%. Our findings in terms of convergence and performances of the OC algorithm are robust with respect to different set of parameters, R_0 and V_s , as tested by increasing/decreasing the initial conditions and changing the velocity amplitude accordingly as, e.g., considering $(10R_0 \rightarrow R_0, 10V_s \rightarrow V_s)$ or $(R_0/10 \rightarrow R_0, V_s/10 \rightarrow V_s)$. A supremacy of the OC strategies is found when looking at the PDF of catching times, t_c , shown in Fig. 2a. Here, we see that OC protocols deplete the probability to search for long times, as shown by comparing the far right tails for $t_c > t_m$, with $t_m \simeq 9\tau_\eta$ the mean capture time for the OC. Furthermore, we have observed that considering the remaining 10% of episodes where OC captures but the heuristics are unsuccessful, we found that OC is also able to capture at much longer timescales, even close to the time horizon. In Fig. 2 we show that the OC protocol provides an advantage also when comparing the mean catching time (Fig. 2b) and an improvement (up to $\times 6$) for the rarest events when compared trajectory-by-trajectory against each of the heuristic reactive strategies (Fig. 2c). Furthermore, all heuristic reactive strategies perform similarly, stressing the intimate limitations of strategies based on local-in-time cues.

Final and intermediate distance statistics. When the target cannot be captured within the allowed time horizon OC is still optimal with regard to the final distance from the target, as imposed by the first term in the RHS of (2), and almost optimal




Fig. 3 Statistics of final distances. PDF of the logarithm of the distances reached by all reactive heuristic strategies, $R_t/R_t^{(OC)}$ at the time, t_t , when the OC trajectories trespass the value 10η . Pure Pursuit (black full rhombus), Surfing Control (blue full triangles) and Perturbative Optimal control (green full squares). The statistic refers to the episodes where OC fails in catching the target.

also at intermediate times. The latter is shown in Fig. 3, where the PDF of the normalized separation, $R_t/R_t^{(OC)}$, for each of the heuristic strategies is shown at the time when the corresponding OC trajectory reaches the separation $R_t^{(OC)} = 10\eta$. Except for a small set of events, the OC solution is always closer to the target as shown by the strong asymmetry between left and right tails.

Conditional growth rate and controllability. To understand the conditions for a successful capture, it is useful to look at the time-evolution of the typical growth of separations along the trajectories of the target and conditioned on the capacity of the OC protocol to succeed/fail in the capture (black full circle in Fig. 4). Conditioning the average growth rate upon capture and non-capture highlights the degree of controllability of the system. Asymptotically, the growth rate is the same for all protocols and given by the Lagrangian Lyapunov exponent λ (dashed straight




Fig. 4 Logarithm of the error growth averaged over the trajectories for all strategies. Pure Pursuit (black full rhombus), Surfing Control (blue full triangles), Perturbative Optimal control (green full squares) and OC (red open circles) are measured only on the unsuccessful episodes (no capture). The growth rate of the uncontrolled searcher (the tracer with black full circles) is shown for both capture and no-capture episodes.

line). For a short time, $t \lesssim 10\tau_\eta$, we see that the two sets behave quite differently. In particular, the unsuccessful episodes are characterized by a growth rate larger than the average (i.e. $> \lambda$) while the successful ones remain close to the average and even slightly below at the very early times. In the same figure, we also show along the unsuccessful episodes the growth of the distance between the target and the (controlled) agent for the heuristic and OC protocols. As one can see, the OC (red open circles) outperforms the reactive strategies at all times, showing that it is not just making the good moves when the final horizon is approaching.

Sensitivity to initial conditions. To test the robustness of the different protocols, we repeated a subset of successful catching episodes with slightly different initial conditions. For the OC strategies we kept the same history of the unperturbed steering protocols, while for the heuristic protocols we allowed to change the reactive control according to the local environment along the perturbed trajectory (see Fig. 5a) for a 3d rendering of a typical numerical experiment comparing OC and PP strategies. Figure 5b displays the PDF of the final distances reached at catching time of the OC unperturbed reference episode by the perturbed OC agent and by the agents following the three heuristic strategies. Results are shown for a subset of the hardest episodes where the unperturbed capture time is much larger than the mean, i.e. $t_c^{(OC)} \gg 9\tau_\eta$, but still all the unperturbed heuristic trajectories capture within the time horizon. OC shows superior robustness against the reactive strategies concerning both the percentage of perturbed trajectories that succeed to capture the target and the final distance from the target for those episodes that do not succeed, as shown by the most pronounced peak at the normalized capture distance, $R/\eta = 10^{-2}$ and by the much shorter right tail developed by the OC episodes with respect to the three reactive protocols (red circles). This result is particularly unexpected because the agents which use heuristic strategies start from perturbed positions but do follow the instantaneous correct control. The performances of the three heuristic strategies would be even worse if we had used, as for the OC agents, the unperturbed strategies learned along the unperturbed trajectory. OC protocols appear to be quite robust, making the target's trajectory an attracting set for the controlled dynamics. This is further quantified

in Fig. 5c where we show for OC and PP strategies the finite time Lyapunov exponent (FTLE) evaluated as a function of the unperturbed capture time, $\gamma_{t_c}^{(OC)} = 1/t_c^{(OC)} \ln(R_{t_c}^{(OC)}/R_0)$ where $R_{t_c}^{(OC)}$ denotes the distance between the perturbed agent and the target^{54–56}. While the characteristic Lyapunov exponent quantifies the average (long term) expansion rate, the FTLE accounts for the (finite-time) fluctuations of the expansion rate, which are a hallmark of intermittent chaotic systems. As a result, the FTLE provides a more quantitative measure of control performance in turbulence. In Fig. 5c the FTLE is an increasing function of $t_c^{(OC)}$, confirming that larger catching times are connected to more chaotic target trajectories. Moreover, FTLE for OC is always smaller than the one for PP, independently of $t_c^{(OC)}$, confirming that OC finds ‘more stable’ trajectories. The findings presented here illustrate the impact of a perturbation in the initial condition of the order of R_c (i.e., $\delta R_0 = R_c = \eta/100$). Anyway, considering other perturbations and limiting them to ensure the velocity field between the pair remains smooth, yield qualitatively similar results.

Discussion and outlook

We have shown how optimal control theory can be efficiently applied to control the dynamics of Lagrangian pairs in turbulent flows. Through the implementation of an iterative algorithm, designed for solving the Euler-Lagrange equations arising from the Pontryagin Minimum Principle^{40,52,53}, we have demonstrated how to catch a moving target in the shortest possible time or, at worst, how to limit the effects of chaotic dispersion.

We focused on the relative dynamics of two Lagrangian objects in 3d turbulence, in the limit where their separation is always smaller than the flow dissipative scale. Both the moving target and the agent are carried by the turbulent flow, but the latter is equipped with some limited propulsion capabilities. Tracking a particle in turbulence is extremely challenging because of the exponential rate of separation due to chaos, making it not obvious a priori that the iterative algorithm converges to an optimal solution. Indeed, each point of an optimal trajectory depends on the entire history (past, present, and future) of the control and, consequently, a high fluctuation of the velocity field at a certain time may affect the entire trajectory, hindering the identification




Fig. 5 Robustness under perturbations of the initial conditions. **a** Example of 100 trajectories obtained by varying the initial conditions of the pursuers. Target reference trajectory (large yellow circles), unperturbed OC trajectory (red open circles), Pure Pursuit strategy (black full rhombus). Lines show the perturbed trajectories, starting with an error $\delta R_0 \approx \eta/10$, the same colors of the unperturbed reference trajectories. **b** PDF of the distances reached at the catching time $t_c^{(OC)}$ by each control strategy starting from perturbed initial conditions (Surfing Control is shown with blue full triangles and Perturbative OC with green full squares). The PDF refers to episodes where OC captures in times larger than the average capture time, i.e., $t_c^{(OC)} \gg t_m^{(OC)}$ and the heuristic strategies capture within the time horizon, t_i . The error on the initial position is set to $\delta R_0 = \eta/100$. The two vertical dashed lines from left to right indicate the capture and the initial distance, respectively. **c** Finite Time Lyapunov Exponent (FTLE) for OC and the Pure Pursuer perturbed strategies evaluated as a function of the unperturbed capture time $t_c^{(OC)}$ for the same episodes of **(b)**. The black full rhombus and the red open circles indicate the mean over all perturbed trajectories, while the shaded areas indicate their standard deviations.

of optimal solutions. Notwithstanding these difficulties, we were able to obtain a high percentage of convergence for the iterative algorithm, paving the way for the application of optimal control theory—so far mostly limited to simple, analytically solvable problems—to turbulent setups as well.

Converging to the optimal solutions requires perfect observability and full knowledge of the system dynamics. One may wonder about the trade-off between hard-to-get but far-sighted optimal solutions and easy-to-apply reactive protocols as, e.g., heuristic controls based only on local and instantaneous information of the environment. In this context, we developed a reactive-control strategy by means of a perturbative approach to solve the optimal control problem. This strategy, which we dubbed *perturbative optimal control*, aims to achieve the best possible performance when operating within the constraint of a short-sighted environmental evaluation. Owing to the assumption of persistent gradient it can be handled analytically. Higher-order expansions of this perturbative approach can in principle be applied, leading to other new heuristic strategies. We found that trading off observability with computational ease inevitably leads to a high loss in performance: more complex control strategies than just reactive ones appear to be necessary. To fill the gap between optimal solutions and heuristic strategies, it would be interesting to improve the latter by introducing short-term memory with some analytical approximation or exploiting data-driven methods.

Remarkably, OC strategies turned out to be highly resilient to disturbances on the agent starting conditions, which highlights how the drifting target trajectory becomes an attractor for the controlled dynamics. Therefore, here the robustness lies in the fact that the controlled trajectory is able to tame the chaotic nature typical of the uncontrolled system. A future perspective would be to investigate how OC and heuristic strategies are impacted by measurement errors (e.g. on position of the Lagrangian pair and on the local velocity gradients) that occur throughout the trajectory, as opposed to errors affecting only the initialization, exploiting tools from stochastic optimal control theory^{57,58}.

The results presented in this study were obtained for a specific set of parameters ($Re_\lambda \approx 130$, $R_0 \approx \eta$, $V_s \approx R_0/\lambda$). While the robustness of the OC algorithm has been tested for different R_0 and V_s parameters, the role of different Re_λ is still an open problem. It would be interesting to explore the importance of adopting optimal navigation strategies when the turbulence intensity varies. In particular, increasing Reynolds number could even lead to an improvement in the relative advantage of OC with respect to heuristic strategies, leveraging on intermittent non-Gaussian fluctuations. This question is left for future work.

As a possible extension of the present results, with an eye to applications such as the control of micro-swimmers at small scales, it could be of interest to include hydrodynamic interactions between particles into the dynamics, as well as to take into

account the impact of a non negligible reorientation time for the control variable.

Finally, an outstanding issue is how to extend the search for moving targets to other situations of interest, for instance when separations are beyond the Kolmogorov scale and the velocity differences are not smooth. Even though heuristic, reactive, and local methods can in principle be employed, gauging their effectiveness is difficult in absence of the benchmark of the optimal control solution. In fact, how to apply optimal control theory in the inertial regime of turbulence remains an open problem, as iteratively solving the Euler-Lagrange equations would require either the simulation of a new 3D Eulerian field at each iteration or the storage of the entire evolution of fields. Both approaches appear to be computationally out of reach, hence the call to develop new numerical techniques to accurately solve large-scale Lagrangian optimization problems.

Methods

Navier-Stokes simulations for Lagrangian tracers. The target trajectories follow the tracer dynamics

$$\dot{\mathbf{X}}_t = \mathbf{u}_t(\mathbf{X}_t), \quad (3)$$

where $\mathbf{u}$ is a solution of the Navier-Stokes equations^{45,46}

$$\begin{cases} \partial_t \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \nu \Delta \mathbf{u} + \mathbf{F} \\ \nabla \cdot \mathbf{u} = 0 \end{cases}, \quad (4)$$

for an incompressible fluid of viscosity ν . The flow is driven to a non-equilibrium statistically steady state by a homogeneous and isotropic forcing, $\mathbf{F}$, obtained via a second-order Ornstein-Uhlenbeck process⁵⁹. For the direct numerical simulations (DNS) we used a standard pseudo-spectral solver fully dealiased with the two-third rule. Details on the simulation can be found in^{60,61}. Parameters of the DNS used in this work are given in Table 1. The database of Lagrangian trajectories used in this study is dumped each $15dt \simeq \tau_\eta/10$.

Optimal control equations. Using the assumption that the distance between the agent and the target is always within the scale of smoothness of the velocity field, i.e. $R_t \lesssim 10\eta$, we can linearize Eq. (1) obtaining

$$\begin{cases} \dot{\mathbf{R}}_t = \nabla \mathbf{u}_t \mathbf{R}_t + \mathbf{U}_t, \\ \mathbf{U}_t = V_s \hat{\mathbf{n}}_t. \end{cases} \quad (5)$$

Here, the pair separation $\mathbf{R}_t$ represents the state variable with a given initial condition $\mathbf{R}_0$ and $\mathbf{U}_t$ is the control variable. As discussed in main text, we aim at solving the following optimization problem: to find the best control that allows the agent to reach the capture distance R_c in minimal time $t_c \leq t_f$, where t_f is a fixed time horizon; if the capture is not realized, we require the control to minimize the final distance, R_{t_f} , from the target. The above twofold goal can be formally imposed by requiring that the

control minimize the following performance index

$$J = R_{t_f}^2 + c\lambda R_c^2 \int_0^{t_f} dt f(R_t), \quad (6)$$

where λ is the uncontrolled Lyapunov exponent and f is an appropriate smooth function (see below) that is equal to 1 for $R_t > R_c$ and 0 for $R_t \leq R_c$. The first term in (6) amounts to requiring minimal distance at time t_f , while the second fulfills the request of minimal time to reach the capture distance. The non-dimensional parameter c weighs the importance of the two objectives. Note that if the capture is not reachable, the second term in the performance index is always a constant, $c\lambda R_c^2 t_f$ and the problem remains minimize the final distance. Conversely, if capture is realized the first term is fixed to R_c^2 . Therefore, to balance the two terms in such a way to favor the capture we must choose $c > 1/(\lambda t_c)$, e.g. using $t_c \sim 1 - 10\tau_\eta$ we can estimate $c \geq 10 - 100$. The results shown here correspond to $c = 100$; while higher values of c lead to same results, $c \sim O(1)$ or even smaller does not allow to find optimal strategies. Given the unconstrained performance index (6), by imposing the dynamics (5) and the non-linear constrain in the control variable, $|\hat{\mathbf{n}}_t|^2 = 1 \forall t$, we are left with the following constrained optimization problem

$$\begin{aligned} \tilde{J} = R_{t_f}^2 + \int_0^{t_f} dt [c\lambda R_c^2 f(R_t) + \\ + \phi_t \cdot (\nabla \mathbf{u}_t \mathbf{R}_t + V_s \hat{\mathbf{n}}_t - \dot{\mathbf{R}}) + \mu_t (1 - |\hat{\mathbf{n}}_t|^2)], \end{aligned} \quad (7)$$

where ϕ_t, μ_t are the Lagrangian multipliers with the role of co-state⁴⁰. Integrating by parts, and requiring the stationarity of $\tilde{J}$ upon variation of control $\hat{\mathbf{n}}_t$ and state $\mathbf{R}_t$ the optimization reduces to solving the following Euler-Lagrange equations⁴⁰:

$$\dot{\phi}_t = -\frac{\partial H_t}{\partial \mathbf{R}_t}, \quad (8)$$

$$\frac{\partial H_t}{\partial \hat{\mathbf{n}}_t} = 0, \quad (9)$$

$H_t = c\lambda R_c^2 f(R_t) + \phi_t \cdot (\nabla \mathbf{u}_t \mathbf{R}_t + V_s \hat{\mathbf{n}}_t) + \mu_t (1 - \hat{\mathbf{n}}_t^2)$ being the Hamiltonian function of the constrained minimization problem. The equation for $\dot{\phi}_t$ has final condition $\phi_{t_f} = 2\mathbf{R}_{t_f}$, while the dynamics of the state variable (5) (which, as discussed below, is actually modified for stopping the dynamics when the capture distance is reached) has initial condition $\mathbf{R}_0$. Notice that Eq. (9) prescribes the control to be

$$\hat{\mathbf{n}}_t = \frac{V_s \phi_t}{2\mu_t} = -\frac{\phi_t}{|\phi_t|}, \quad (10)$$

where μ_t plays the role of normalization factor and where the minus sign is to impose $\frac{\partial^2 H_t}{\partial \hat{\mathbf{n}}_t^2} > 0$ as we are performing a minimization. As a result, the attainment of optimality requires to satisfy a boundary conditions problem. This implies that the control that is optimal at a certain time is inherently linked to the past and future evolution of the trajectory, rather than solely relying on the local flow in the vicinity of the current position of the agent. In other words the control strategy is not optimal at each time step but globally. As a consequence, minimizing the capture time cannot be achieved by a piecewise trajectories optimization, i.e. considering shorter time horizons. Indeed, from eq.(6) the boundary conditions would impose the minimization of the distance at time t_f , without taking into account the fact that to capture the target in the future (e.g., in the next time horizon) there might have been a more convenient path to follow in the past.

Table 1 Parameters of the DNS: N resolution in each dimension; L physical dimension of the 3-periodic box; dt time step in the DNS integration; ν kinematic viscosity; $\epsilon = \nu \langle \partial_i \mu_j \partial_i \mu_j \rangle$; $\tau_\eta = \sqrt{\nu/\epsilon}$; $\eta = (\nu^3/\epsilon)^{1/4}$; $Re_\lambda = u_{rms}\lambda/\nu$, where $\lambda = \sqrt{5E_{tot}/\Omega}$ is the 'Taylor-scale' measured from the ratio between the mean system energy and enstrophy.

N	L	dt	ν
1024	2π	1.5×10^{-4}	8×10^{-4}
ϵ	τ_η	η	Re_λ
1.4 ± 0.1	0.023 ± 0.003	0.0042 ± 0.0001	$\simeq 130$

Table 2 Pseudocode of the optimal control algorithm.**Algorithm:** Forward-Backward SweepParameters: learning rate γ , threshold for convergence δ 1: Initial guess for the control variable $\hat{\mathbf{n}}_t \forall t \in [0, t_f]$ 2: Forward integration of $\mathbf{R}_t$ in $t \in [0, t_f]$ (Eq. (5))3: Backward integration of ϕ_t in $t \in [t_f, 0]$ (Eq. (8))4: Update control $\hat{\mathbf{n}}_t \leftarrow (1 - \gamma)\hat{\mathbf{n}}_t + \gamma \frac{\phi_t}{|\phi_t|}$ (Eq. (10))5: Check convergence **if** $\sum_{i=1}^{t_f/dt} |\Delta w_{idt}| \leq \delta \sum_{i=0}^{t_f/dt} |w_{idt}|$
for all $\mathbf{w} \in \{\phi, \mathbf{R}, \hat{\mathbf{n}}\}$ **end else** goto 2where Δw_i is the difference between the old and new estimate of the variables
Forward-backward sweep method for iterative solution of the Euler-Lagrange equations.

In practice, to find the optimal control solution that minimizes $\tilde{J}$, we use the Forward-Backward Sweep Method (FBSM)⁵². At first, it requires an initial guess for the control variable $\hat{\mathbf{n}}_t$ for $0 \leq t \leq t_f$. Then, the problem becomes computationally heavy since it requires iterative searching with backward (for the Lagrangian multipliers) and forward (for the state variable) integration such as to identify the optimal control. The algorithm can be summarized with the following pseudocode, Table 2.

In the practical implementation we used $\gamma = 5 \times 10^{-4}$, $\delta = 10^{-4}$ and as initial guess for the control variable we used the best heuristic strategy, i.e., the one with the minimum performance index. Using PP for the initialization provides quantitatively similar results. If the convergence is not realized after 20K iterations of the algorithm, we repeated the FBSM rescaling the parameters $\gamma < \gamma/10$, $\delta < \delta/10$ and increasing the maximum number of iteration to 2×10^5 . This setup ensure the convergence in the 81% of the total episodes studied. In the other 19% of episodes, the FBSM algorithm does not converge, so the global optimum is not achieved. However, it is always possible to consider the strategy that during the iterations of the FBSM algorithm provides the minimum of the performance index as an approximate solution of the optimal control. Indeed, since the initialization of the algorithm is provided by the best of the heuristic strategies, the approximate solution will always be an upper bound of the reactive behavior. By refining the hyperparameters even more, should be possible to make the algorithm to always converge. Moreover for the function f in Eq. (6) we used

$$f = \frac{1}{2} \tanh\left(\alpha \frac{R_t - R_c}{R_c}\right) \quad (11)$$

which is a smoothed version of the Heaviside function, whose stiffness is ruled by α (here chosen to be 10). In addition, to impose the capture condition, i.e. that whenever the agent is at distance $R \leq R_c$ from the target it sticks to it we modified Eq. (5) such that $\dot{\mathbf{R}}_t \rightarrow 0$ when $R_t \leq R_c$ by redefining the dynamics as follows:

$$f \dot{\mathbf{R}}_t \rightarrow \dot{\mathbf{R}}_t, \quad (12)$$

which implies that the backward evolution of the Lagrangian multipliers in (8) becomes

$$\dot{\phi}_t = -[c\lambda R_c^2 + (\nabla \mathbf{u}_t \mathbf{R}_t + V_s \hat{\mathbf{n}}_t) \cdot \phi_t] \frac{\partial f}{\partial \mathbf{R}_t} - f \nabla \mathbf{u}_t^\top \phi_t. \quad (13)$$

Heuristic strategies. Differently from the optimal control equations, the heuristic control strategies we studied are reactive, meaning that they need only instantaneous information about the system to be applied. The pure pursuit (PP) strategy, for instance, does not exploit any information on the flow during the navigation but constantly realign the control direction to the moving target, i.e. $\hat{\mathbf{n}}_t^{\text{PP}} = -\dot{\mathbf{R}}_t$. The other two strategies, the surfing

control (SC) and the perturbative optimal control (PO), instead, consider both the direction of the target and the instantaneous velocity gradient. While these two strategies are both based on a free parameter to be optimized numerically, τ_s and τ_p respectively, they are obtained from different assumptions.

Surfing control strategy. This control is inspired by Ref. 51, where one assume that a linear approximation of the flow underlying the active particle is reasonable for an interval of time τ_s . In other words, one has to assume τ_s as the persistence time of the local (to the agent) gradients of the flow. This control was proposed as an effective strategy to drift in a given direction (e.g. the vertical one) fixed in time, essentially it is obtained with the request to maximize over the time interval τ_s the displacement in the chosen direction. Here, we adapt it to our case by assuming that the direction toward the target, $\hat{\mathbf{R}}_t$, remains constant during the same interval of time τ_s . Then, maximizing the searcher displacement along the target direction, i.e. $\max_{\hat{\mathbf{n}}_t} [-(\mathbf{X}_{t+\tau_s}^{(a)} - \mathbf{X}_t^{(a)}) \cdot \hat{\mathbf{R}}_t]$ and considering a continuous measurement of the environment, we find the following strategy:

$$\hat{\mathbf{n}}_t^{\text{SC}} = -\frac{[\exp(\tau_s \nabla \mathbf{u}_t)]^\top \hat{\mathbf{R}}_t}{|[\exp(\tau_s \nabla \mathbf{u}_t)]^\top \hat{\mathbf{R}}_t|}. \quad (14)$$

The details on the derivation are discussed in Ref. 51 and are not repeated here. However, it is important to note the role of the exponential in the control variable, as demonstrated in⁵⁰. In particular, by expanding in series eq.(14) it is possible to obtain

$$\begin{aligned} \mathbf{n}_t^{(\text{SC})} &\propto -\left[\sum_{k=0}^{\infty} \frac{\tau_s^k}{k!} ((\nabla \mathbf{u}_t)^\top)^k\right] \mathbf{R}_t \\ &= -\left[\mathbf{R}_t + \tau_s \nabla \mathbf{u}_t \mathbf{R}_t + \frac{1}{2} \tau_s^2 \nabla(\mathbf{u}_t \nabla \mathbf{u}_t \mathbf{R}_t) + \dots\right] = \sum_{k=0}^{\infty} \mathbf{n}_{t,k}, \end{aligned} \quad (15)$$

with $\mathbf{n}_{t,0} = -\mathbf{R}_t$ and $\forall k > 0$ $\mathbf{n}_{t,k} = \frac{1}{k} \tau_s [\nabla \mathbf{u}_t] \mathbf{n}_{t,k-1}$. This means that surfing control is achieved through a weighted sum of several directions, $\mathbf{n}_{t,k}$, where the weight is defined by τ_s . These directions are obtained as the gradient of the local field, $\mathbf{u}_t$, projected along the direction of the previous term in the series, $\mathbf{n}_{t,k-1}$. This allows the surfing strategy to reorient the order 0 of the series (Pure Pursuit strategy) under the influence of beneficial currents.

Perturbative optimal control strategy. The PO strategy we propose, works in the same regime we derived optimal control equations, considering the velocity field between the agent and the target to be linearizable (5) but, similarly to the surfing, by assuming local gradients persistence for a time τ_p . Under these assumptions, it is easy to show that

$$\mathbf{R}_{\tau_p} = \exp(\tau_p \nabla \mathbf{u}_0) \mathbf{R}_0 + V_s \int_0^{\tau_p} dt \left[\exp[(\tau_p - t) \nabla \mathbf{u}_0] \hat{\mathbf{n}}_t^{\text{PO}} \right]. \quad (16)$$

Then, we derive the PO strategy by imposing that $\mathbf{R}_{\tau_p} \cdot \mathbf{R}_{\tau_p}^{V_s=0}$ is minimum (note that this perturbative approach can be thought as the 0th order, in V_s , solution of the optimal control equations), i.e., minimizing directly the separation of the pair. Finally, assuming continuous measurements as before, we obtain the following control:

$$\hat{\mathbf{n}}_t^{\text{PO}} = -\frac{[\exp(\tau_p \nabla \mathbf{u}_t)]^\top \exp(\tau_p \nabla \mathbf{u}_t) \hat{\mathbf{R}}_t}{|[\exp(\tau_p \nabla \mathbf{u}_t)]^\top \exp(\tau_p \nabla \mathbf{u}_t) \hat{\mathbf{R}}_t|}. \quad (17)$$

Note that SC and PO strategies recover PP for the free parameters, τ_s and τ_p , set to be zero. Clearly, the optimal free




Fig. 6 Control strategies in a simple 2d flow. Trajectories obtained in a 2-dimensional steady linear flow. Pure Pursuit strategy (black), Surfing Control (blue), Perturbative Optimal Control (green) and Optimal Control (red). The trajectories are shown in the relative coordinate space, $\mathbf{R} = (R_x, R_y)$, thus the target is fixed in the origin (full yellow circle). The dashed lines show the trajectories followed by the strategies after the capture time of OC. The arrows indicate the control directions followed by each strategy.

parameters depend on the temporal variation of both the underlying gradients, $\nabla \mathbf{u}_p$, and the target direction, $\hat{\mathbf{R}}_t$. As a result, it is tempting to guess that they will be proportional to τ_η . Their values have been determined empirically, searching for the values that maximize the capture frequencies. In particular, we have found $\tau_s \simeq 0.6\tau_\eta$ and $\tau_p \simeq 1.3\tau_\eta$.

To get an intuition of the difference between the strategies we show them in simple steady flows. Consider for simplicity a 2-dimensional case, where the target is fixed at the origin $(0, 0)$ in the coordinate space $\mathbf{R} = (R_x, R_y)$. In a hyperbolic field, it is easy to see that SC and PO strategies are equivalent as long as we define τ_p as half of τ_s . Additionally, unlike PP, these strategies leverage the local gradient to reach the contracting axis of the field which directly leads them towards the target. In an elliptic field, PO and PP strategies are identical and optimal. Instead, the SC is not optimal and approaches the target while counter-rotating with respect to the direction induced by the flow.

As an illustration, in Fig. 6 we show the trajectories and controls followed by the strategies in a generic steady field. SC and PO utilize the local gradient to navigate towards regions where the flow is weaker and the target can be reached more efficiently. Furthermore, in this case the PO strategy nearly performs as good as the OC.

Data availability

The Lagrangian target trajectories (including positions, velocities, accelerations and fluid gradients along each particle) used in this work are available for download in the Smart-TURB portal <http://smart-turb.roma2.infn.it>, under the TURB-Lagr repository⁶⁰. TURB-Lagr is a new open database of 3d turbulent Lagrangian trajectories, obtained by Direct Numerical Simulations (DNS) of the Navier-Stokes equations with homogeneous and isotropic forcing. Details on how to download and read the database are also given in the

portal. The analysis that support the findings of this study are available from the corresponding author upon reasonable request.

Code availability

The code (written in C language) to study the optimal tracking strategies in a turbulent flow is free downloadable on GitHub at this link: <https://github.com/SmartTURB/Optimal-tracking-strategies-in-a-turbulent-flow>.

Received: 2 May 2023; Accepted: 31 August 2023;

Published online: 29 September 2023

References

- Trincavelli, M. et al. Towards environmental monitoring with mobile robots. In *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2210–2215 (IEEE, 2008).
- Zhang, W., Inanc, T., Ober-Blobaum, S. & Marsden, J. E. Optimal trajectory generation for a glider in time-varying 2d ocean flows b-spline model. In *2008 IEEE International Conference on Robotics and Automation*, 1083–1088 (IEEE, 2008).
- Bellemare, M. G. et al. Autonomous navigation of stratospheric balloons using reinforcement learning. *Nature* **588**, 77–82 (2020).
- Chai, F. et al. Monitoring ocean biogeochemistry with autonomous platforms. *Nat. Rev. Earth Environ.* **1**, 315–326 (2020).
- Wang, J. & Gao, W. Nano/microscale motors: biomedical opportunities and challenges. *ACS Nano* **6**, 5745–5751 (2012).
- Li, J., Esteban-Fernández de Ávila, B., Gao, W., Zhang, L. & Wang, J. Micro/nanorobots for biomedicine: Delivery, surgery, sensing, and detoxification. *Sci. Rob.* **2**, eaam6431 (2017).
- Wang, B., Kostarelos, K., Nelson, B. J. & Zhang, L. Trends in micro-/nanorobotics: materials development, actuation, localization, and system integration for biomedical applications. *Adv. Mater.* **33**, 2002047 (2021).
- Szczerba, R. J., Galkowski, P., Glickstein, I. S. & Ternullo, N. Robust algorithm for real-time route planning. *IEEE Trans. Aerospace Electr. Syst.* **36**, 869–878 (2000).
- Song, Z., Lipinski, D. & Mohseni, K. Multi-vehicle cooperation and nearly fuel-optimal flock guidance in strong background flows. *Ocean Eng.* **141**, 388–404 (2017).
- Guerrero, J. & Bestaoui, Y. Uav path planning for structure inspection in windy environments. *J. Intell. Robot. Syst.* **69**, 297–311 (2013).
- Nasiri, M., Löwen, H. & Liebchen, B. Optimal active particle navigation meets machine learning. *Europhys. Lett.* **142**, 17001 (2023).
- Lolla, T., Lermusiaux, P. F. J., Ueckermann, M. P. & Haley, P. J. Time-optimal path planning in dynamic flows using level set equations: theory and schemes. *Ocean Dyn.* **64**, 1373–1397 (2014).
- Rhoads, B., Mezić, I. & Poje, A. C. Minimum time heading control of underpowered vehicles in time-varying ocean currents. *Ocean Eng.* **66**, 12–31 (2013).
- Biferale, L., Bonaccorso, F., Buzzicotti, M., Clark Di Leoni, P. & Gustavsson, K. Zermelo's problem: optimal point-to-point navigation in 2d turbulent flows using reinforcement learning. *Chaos: Interdiscip. J. Nonlinear Sci.* **29**, 103138 (2019).
- Buzzicotti, M., Biferale, L., Bonaccorso, F., Clark di Leoni, P. & Gustavsson, K. Optimal control of point-to-point navigation in turbulent time dependent flows using reinforcement learning. In *AIxLA 2020 – Advances in Artificial Intelligence*, 223–234 (Springer International Publishing, Cham, 2021).
- Alageshan, J. K., Verma, A. K., Bec, J. & Pandit, R. Machine learning strategies for path-planning microswimmers in turbulent flows. *Phys. Rev. E* **101**, 043110 (2020).
- Daddi-Moussa-Ider, A., Löwen, H. & Liebchen, B. Hydrodynamics can determine the optimal route for microswimmer navigation. *Commun. Phys.* **4**, 15 (2021).
- Gunnarson, P., Mandralis, I., Novati, G., Koumoutsakos, P. & Dabiri, J. O. Learning efficient navigation in vortical flow fields. *Nat. Commun.* **12**, 7143 (2021).
- Verma, S., Novati, G. & Koumoutsakos, P. Efficient collective swimming by harnessing vortices through deep reinforcement learning. *Proc. Natl. Acad. Sci.* **115**, 5849–5854 (2018).
- Goh, S., Winkler, R. G. & Gompper, G. Noisy pursuit and pattern formation of self-steering active particles. *New J. Phys.* **24**, 093039 (2022).
- Zhu, G., Fang, W.-Z. & Zhu, L. Optimizing low-reynolds-number predation via optimal control and reinforcement learning. *J. Fluid Mech.* **944**, A3 (2022).
- Yang, Y. & Bevan, M. A. Optimal navigation of self-propelled colloids. *ACS Nano* **12**, 10712–10724 (2018).

23. Yang, Y., Bevan, M. A. & Li, B. Efficient navigation of colloidal robots in an unknown environment via deep reinforcement learning. *Adv. Intell. Syst.* **2**, 1900106 (2020).
24. Piro, L., Mahault, B. & Golestanian, R. Optimal navigation of microswimmers in complex and noisy environments. *New J. Phys.* **24**, 093037 (2022).
25. Piro, L., Golestanian, R. & Mahault, B. Efficiency of navigation strategies for active particles in rugged landscapes. *Front. Phys.* **10**, 1125 (2022).
26. Calascibetta, C., Biferale, L., Borra, F., Celani, A. & Cencini, M. Taming lagrangian chaos with multi-objective reinforcement learning. *Eur. Phys. J. E* **46**, 9 (2023).
27. Xu, A., Wu, H.-L. & Xi, H.-D. Long-distance migration with minimal energy consumption in a thermal turbulent environment. *Phys. Rev. Fluids* **8**, 023502 (2023).
28. Peterson, C. & Paley, D. Multivehicle coordination in an estimated time-varying flowfield. *J. Guid. Control Dyn.* **34**, 177–191 (2011).
29. Song, Z. & Mohseni, K. Anisotropic active lagrangian particle swarm control in a meandering jet. In *2015 54th IEEE Conference on Decision and Control (CDC)*, 240–245 (IEEE, 2015).
30. Mallory, K., Hsieh, M., Forgoon, E. & Schwartz, I. Distributed allocation of mobile sensing swarms in gyre flows. *Nonlinear Processes Geophys.* **20**, 657–668 (2013).
31. Wynn, R. B. et al. Autonomous underwater vehicles (auvs): Their past, present and future contributions to the advancement of marine geoscience. *Mar. Geol.* **352**, 451–468 (2014).
32. Witt, J. & Dunbabin, M. Go with the flow: Optimal auv path planning in coastal environments. *Proceedings of the 2008 Australasian Conference on Robotics and Automation, ACRA 2008* (2008).
33. Smith, R., Das, J., Hine, G., Anderson, W. & Sukhatme, G. Predicting wave glider speed from environmental measurements. *OCEANS'11 - MTS/IEEE Kona, Program Book* (2011).
34. Lumpkin, R. & Pazos, M. *Measuring surface currents with Surface Velocity Program drifters: the instrument, its data, and some recent results*, 39–67 (Cambridge University Press, 2007).
35. Bechinger, C. et al. Active particles in complex and crowded environments. *Rev. Mod. Phys.* **88**, 045006 (2016).
36. Kurzthaler, C. et al. Probing the spatiotemporal dynamics of catalytic janus particles with single-particle tracking and differential dynamic microscopy. *Phys. Rev. Lett.* **121**, 078001 (2018).
37. Popescu, M. N., Tasinkevych, M. & Dietrich, S. Pulling and pushing a cargo with a catalytically active carrier. *Europhys. Lett.* **95**, 28004 (2011).
38. Baraban, L. et al. Transport of cargo by catalytic janus micro-motors. *Soft Matter* **8**, 48–52 (2012).
39. Panda, M., Das, B., Subudhi, B. & Pati, B. B. A comprehensive review of path planning algorithms for autonomous underwater vehicles. *Int. J. Autom. Comput.* **17**, 321–352 (2020).
40. Bryson, A. E. *Applied Optimal Control: Optimization, Estimation and Control (1st ed.)* (Routledge, 1975).
41. Ben-Asher, J. Z. *Optimal Control Theory with Aerospace Applications* (American Institute of Aeronautics and Astronautics, 2010).
42. Liebchen, B. & Löwen, H. Optimal navigation strategies for active particles. *Europhys. Lett.* **127**, 34003 (2019).
43. Hays, G. et al. Route optimisation and solving zermelo's navigation problem during long distance migration in cross flows. *Ecol. Lett.* **17**, 137–143 (2013).
44. Nahin, P. J. *Chases and escapes: the mathematics of pursuit and evasion* (Princeton University Press, 2012).
45. Frisch, U. *Turbulence: the legacy of AN Kolmogorov* (Cambridge University Press, 1995).
46. Pope, S. B. *Turbulent Flows* (Cambridge University Press, 2000).
47. Cencini, M., Cecconi, F. & Vulpiani, A. *Chaos: From Simple Models to Complex Systems*. Series on advances in statistical mechanics (World Scientific, 2010).
48. Saw, E. W., Shaw, R. A., Ayyalasomayajula, S., Chuang, P. Y. & Gylfason, A. Inertial clustering of particles in high-reynolds-number turbulence. *Phys. Rev. Lett.* **100**, 214501 (2008).
49. Ishihara, T., Gotoh, T. & Kaneda, Y. Study of high-reynolds number isotropic turbulence by direct numerical simulation. *Ann. Rev. Fluid Mech.* **41**, 165–180 (2009).
50. Monthiller, R. A. *mechanistic approach to plankton migration*. PhD dissertation, Central Méditerranée https://github.com/rmonthil-phd/thesis-a-mechanistic-approach-to-plakton-migration/releases/download/v1.0-comments-addressed/thesis_a_mechanistic_approach_to_plankton_migration_v1.pdf (2022).
51. Monthiller, R., Loisy, A., Koehl, M. A., Favier, B. & Eloy, C. Surfing on turbulence: a strategy for planktonic navigation. *Phys. Rev. Lett.* **129**, 064502 (2022).
52. Lenhart, S. & Workman, J. T. *Optimal Control Applied to Biological Models*. Chapman & Hall/CRC Mathematical and Computational Biology (Taylor & Francis, 2007).
53. Trélat, E. Optimal control and applications to aerospace: some results and challenges. *J. Optim. Theory Appl.* **154**, 713–758 (2012).
54. Brunton, S. L. & Rowley, C. W. Fast computation of finite-time Lyapunov exponent fields for unsteady flows. *Chaos: Interdiscip. J. Nonlinear Sci.* **20**, 017503 (2010).
55. Krishna, K., Song, Z. & Brunton, S. L. Finite-horizon, energy-efficient trajectories in unsteady flows. *Proc. R. Soc. A* **478**, 20210255 (2022).
56. Krishna, K., Brunton, S. L. & Song, Z. Finite time lyapunov exponent analysis of model predictive control and reinforcement learning. *arXiv preprint arXiv:2304.03326* (2023).
57. Fleming, W. H. & Rishel, R. W. *Deterministic and stochastic optimal control*, vol. 1 (Springer Science & Business Media, 2012).
58. Crespo, L. G. & Sun, J.-Q. Stochastic optimal control via bellman's principle. *Automatica* **39**, 2109–2114 (2003).
59. Sawford, B. L. Reynolds number effects in Lagrangian stochastic models of turbulent dispersion. *Phys. Fluids A: Fluid Dyn.* **3**, 1577–1586 (1991).
60. Biferale, L., Bonaccorso, F., Buzzicotti, M. & Calascibetta, C. Turb-lagr. a database of 3d lagrangian trajectories in homogeneous and isotropic turbulence. *arXiv preprint arXiv:2303.08662* (2023).
61. Buzzicotti, M., Bhatnagar, A., Biferale, L., Lanotte, A. S. & Ray, S. S. Lagrangian statistics for navier-stokes turbulence under fourier-mode reduction: Fractal and homogeneous decimations. *New J. Phys.* **18**, 113047 (2016).

Acknowledgements

This work was supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant Agreement No. 882340).

Author contributions

C.C., L.B., F.B., A.C., and M.C. conceived the research. C.C. performed all the numerical simulations and data analysis. C.C., L.B., F.B., A.C., and M.C. contributed to the interpretation of the results and writing of the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s42005-023-01366-y>.

Correspondence and requests for materials should be addressed to Chiara Calascibetta.

Peer review information : *Communications Physics* thanks Abdallah Daddi-Moussa-Ider, and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. A peer review file is available.

Reprints and permission information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023