

Energy-positive soaring using transient turbulent fluctuations

Danyun He,^{1,2,*} Gautam Reddy,^{3,4,5} and Chris H. Rycroft^{2,6,†}

¹*John A. Paulson School of Engineering and Applied Sciences,
Harvard University, Cambridge, MA 02138, USA*

²*Department of Mathematics, University of Wisconsin–Madison, Madison, WI 53706, USA*

³*Physics & Informatics Laboratories, NTT Research, Inc., Sunnyvale, CA 94085, USA*

⁴*Center for Brain Science, Harvard University, Cambridge, MA 02138, USA*

⁵*NSF-Simons Center for Mathematical & Statistical Analysis of Biology,
Harvard University, Cambridge, MA 02138, USA*

⁶*Computational Research Division, Lawrence Berkeley Laboratory, Berkeley, CA 94720, USA*

(Dated: April 13, 2023)

Soaring birds gain energy from stable ascending currents or shear. However, it remains unclear whether energy loss due to drag can be overcome by extracting work from transient turbulent fluctuations. We designed numerical simulations of gliders navigating in a kinematic model that captures the spatio-temporal correlations of atmospheric turbulence. Energy extraction is enabled by an adaptive algorithm based on Monte Carlo tree search that dynamically filters acquired information about the flow to plan future paths. We show that net energy gain is feasible under realistic constraints. Glider paths reflect patterns of foraging, where exploration of the flow is interspersed with bouts of energy extraction through localized spirals.

Soaring birds harvest energy by strategically gliding through atmospheric flows while extracting energy contained in the flows. For soaring birds like herring gulls and albatrosses, gliding consumes oxygen at a rate $\approx 30\%$ lower than flapping [1, 2], which is crucial for making long-distance migration feasible within metabolic and aerodynamic constraints [3]. The energetics of different forms of soaring can be described by a general expression for the non-dimensionalized rate of energy gained, $\dot{\epsilon}$, by a glider (that generates no thrust) in a wind field:

$$\dot{\epsilon} = -c_D v^3 + w_z - \mathbf{v} \cdot \dot{\mathbf{w}}, \quad (1)$$

where gravity points in the negative z direction, c_D is the aerodynamic drag coefficient, $v = \|\mathbf{v}\|$ is the airspeed of the glider and $\mathbf{w} = (w_x, w_y, w_z)$ is the wind velocity [4–6]. (See Supplemental Material (SM) section SM.II for details.) In the absence of wind, energy is continuously lost due to drag and the glider sinks at a constant rate. Thus, in order to compensate for drag, gliders should either actively localize at updrafts or align themselves anti-parallel to rapid gusts and wind shear along the glider’s trajectory. These two mechanisms correspond to the second and third terms in Eq. (1) respectively. Birds exploit both these mechanisms to gain energy, which correspond to two commonly observed modes of soaring known as thermal and dynamic soaring.

Thermal soaring relies on ascending currents (thermals) generated by convection in the atmospheric boundary layer [7–9]. Thermals are dynamic flow structures that typically last a few minutes, providing updrafts that enable a bird to spiral up the boundary layer and forage for prey or glide to another thermal during migration [10]. Adult vultures have more control on centering thermals than juveniles, suggesting that vultures improve their soaring skills through experience [11]. Dynamic soaring allows

for energy-neutral flight over oceans, where thermals are weak or absent [12–14]. The predominant contribution to energy extraction during dynamic soaring is through the third term in Eq. (1) [15], that is, by maintaining an appropriate heading while the bird manoeuvres through a stable shear layer generated behind ocean waves.

Nevertheless, both thermal and dynamic soaring rely on the formation of relatively stable convective plumes or wind shear. Atmospheric turbulent flows contain short-lived eddies of multiple time scales with velocity amplitudes comparable to a glider’s typical sink rate, offering a potential continuous source of energy even in environments without stable convection or shear [16, 17]. It is unclear whether energy can be extracted solely from the rapidly fluctuating wind fields that characterize such turbulent flows [18, 19]. In order to harvest energy effectively, a glider should first acquire information about the flow, identify potential energy sources and select its path to maximize energy gain. However, acquired information degrades with time in an unsteady flow, which imposes strong constraints. In the extreme case of an uncorrelated flow, the value of information degrades immediately; consequently, $\dot{\epsilon}$ is negative on average, $\langle \dot{\epsilon} \rangle \simeq -c_D v^3$. This is no longer true when the flow has correlations longer or comparable to the typical control timescale of the glider. What then is the relationship between the maximal energy rate, $\langle \dot{\epsilon} \rangle_{\max}$, the dynamic properties of the flow and the aerodynamic constraints on a glider? Even if energy-positive soaring is physically plausible, i.e., $\langle \dot{\epsilon} \rangle_{\max} > 0$, can a computational algorithm feasibly attain this limit?

Previous studies have examined navigational strategies for dynamic soaring in a static shear layer [13, 14] and thermal soaring in convective flows with [20, 21] and without turbulence [22–24]. In the former case, a classical result from Rayleigh establishes the environmental condi-

tions required to achieve energy-neutral flight [25]. For thermal soaring, reinforcement learning methods have proved fruitful for identifying useful cues and effective navigational strategies in the face of turbulence [20, 21]. These settings consider soaring in a convective flow, where the ascending branch is on the scale of hundreds of meters and lasts for minutes. Here, turbulence plays a disruptive role by introducing ‘noise’ when identifying and localizing within relatively stable thermals [26, 27]. A recent proposal to extract energy from unsteady flows relies on a phenomenon known as fast-tracking exhibited by Stokesian inertial particles [28]. Online algorithms that learn the flow in real time to adaptively plan subsequent paths, and more generally, methods for active navigation in complex flows remain unexplored. In this Letter, we address these questions using a kinematic model of three-dimensional homogeneous and isotropic turbulence. We consider a glider navigating within this flow that orients itself in response to its sensory history using a general-purpose decision-making algorithm that combines statistical inference and long-term planning. We show that energy-positive soaring is feasible and delineate the aerodynamic and flow parameters where this can be achieved.

Model. We consider a glider navigating in a fluctuating flow by controlling its bank angle μ while making local wind velocity measurements \mathbf{w} . The equations of the motion of the glider are given by $\dot{\mathbf{u}} = \mathbf{L} + \mathbf{D} + m\mathbf{g}$, where $\mathbf{u} = \mathbf{v} + \mathbf{w}$ is the glider’s velocity with respect to the ground and \mathbf{L}, \mathbf{D} are the lift and drag forces on the glider respectively (see SM.I). Based on the glider’s airspeed in steady flight, we define the characteristic speed and control timescales of the glider as $v_c = \sqrt{2mg/\rho S}$ and $t_c = v_c/g$ respectively, where ρ is the density of air and S is the surface area of the wings. For soaring avian




FIG. 1. The numerical simulation of the three-dimensional homogeneous and isotropic turbulence. The figure shows vertical wind velocity field. Red and blue indicates large ascending and descending currents respectively.

migrants, the typical gliding speed is 14 m/s [29, 30]; accordingly, in what follows we set $v_c = 10$ m/s and the typical control timescale $t_c = 1$ s.

Next, we consider a kinematic model that captures the broad range of length scales and timescales characteristic of turbulent flows. The three-dimensional wind field, $\mathbf{w}(\mathbf{r}, t)$, at position \mathbf{r} and time t is specified in terms of its Fourier modes [31], which have stochastic dynamics with energy spectrum $E(k) \sim k^{-5/3}$ between the smallest and largest wavelength modes k_{\max} and k_{\min} (see SM.III). We use a cubic domain with side length 500 m for the fluid field, with periodic boundary conditions. The fluid is simulated using 64^3 modes. Figure 1 shows an example of the simulated wind field. The root-mean-square (RMS) wind speed is set to be 0.5 m/s based on typical amplitudes of atmospheric wind fluctuations. The modes decay with a characteristic timescale $\tau(k) = \tau_0 k^{-2/3}$, with a prefactor τ_0 that sets the overall scale of temporal fluctuations relative to t_c . The statistical properties of the flow are determined by the stationary, isotropic covariance kernel of each wind velocity component,

$$K(r, t) = \alpha \int_{k_{\min}}^{k_{\max}} k^{-5/3} \text{sinc}(kr) e^{-t/\tau(k)} dk, \quad (2)$$

where $r = \|\mathbf{r}\|$ and α is a constant that sets the RMS wind speed (see SM.III).

We now describe the navigational strategy, which maps the glider’s wind velocity and positional history to one of three bank angle changes, namely, $\Delta\mu = 0^\circ, \pm 10^\circ$ to a maximum (minimum) of $\mu = \pm 40^\circ$. A glider with memory size M and history H_t at time t contains the glider’s past locations and wind measurements collected at fixed intervals $\Delta t = t_c/2$, that is, $H_t = \mathbf{r}_{t-M\Delta t} \mathbf{w}_{t-M\Delta t} \mathbf{r}_{t-(M-1)\Delta t} \mathbf{w}_{t-(M-1)\Delta t} \dots \mathbf{r}_t \mathbf{w}_t$. The glider chooses an action so as to maximize the total expected energy gained in time $d\Delta t$, $\langle \varepsilon_{t+d\Delta t} \rangle_{H_t} - \varepsilon_t$, where $d\Delta t = 10t_c$ sets the planning horizon (see SM.VI for a discussion of this choice). This maximal expected energy gained given a particular history, $V(H_t)$, satisfies the recursive Hamilton–Jacobi–Bellman equation [32],

$$V(H_t) = \max_{a \in \mathcal{A}} \langle \varepsilon_{t+\Delta t} - \varepsilon_t + V(H_{t+\Delta t}) \rangle_{H_t, a}, \quad (3)$$

where the expectation is over wind configurations encountered along the glider’s trajectory from t to $t + \Delta t$ with boundary condition $V(H_{t+d\Delta t}) = 0$. The optimal action is the one that attains the maximum over the set \mathcal{A} of all possible actions that the glider can take. The expectation implicitly contains the ‘propagator’, i.e., the conditional probability density of the wind velocity at a new location, $P(\mathbf{w}_{t'} | H_t, \mathbf{r}_{t'})$ for $t' > t$. Computing the optimal action is generally challenging due to the evaluation of the exponential number ($\sim |\mathcal{A}|^d$) of possible future paths and the wind configurations encountered along these paths. We now describe three simplifications made to derive an

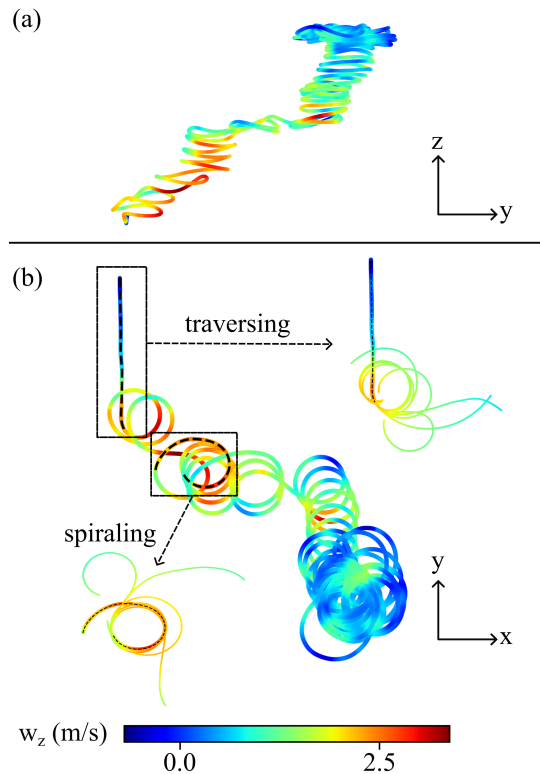


FIG. 2. The glider’s trajectory in a frozen wind field. The color shows the vertical wind velocity that applied on the glider. (a) The glider’s trajectory in yz -plane. (b) The glider’s trajectory in xy -plane. The glider exhibits ‘traversing’ behavior at the beginning of the searching and a switch to ‘spiraling’ behavior once an ascending current has been found. The zoomed-in figures show some examples of the MCTS paths.

efficient, online algorithm from the general expression Eq. (3) (see SM.IV & SM.V for full details).

First, we observe that the glider’s future trajectory and energy gain have approximately a linear dependence on \mathbf{w} if w is much smaller than the typical airspeed of the glider ($w \ll v_c$). Since the flow is defined by its second-order statistics in Eq. (2), it is a Gaussian process. We perform Gaussian process regression to predict wind velocities in future time and locations with given history H_t , and then we compute the expected energy gain along a large number of paths (here, 10^4 paths).

Second, we numerically evaluate the optimal path using Monte Carlo sampling. However, the large number of future paths remains a challenge and efficient pruning techniques are required to make planning tractable. To prune sub-optimal paths during planning, we implement a Monte Carlo tree search (MCTS) [33]. MCTS uses a tree search algorithm to balance exploration and exploitation of future paths. The algorithm stochastically samples different paths, and at each branch chooses an action based on the expected energy gain. The sampling is biased toward paths with higher energy gain, but with

an additional exploration bonus to ensure less promising paths still have a chance of being explored.

Finally, we employ MCTS to find the sequence of actions a_1, a_2, \dots, a_d that approximately maximize the expected energy gain, i.e., $V(H_t) \approx \max_{a_1, a_2, \dots, a_d} (\varepsilon_{t+d\Delta t} - \varepsilon_t)_{H_t, a_1, a_2, \dots, a_d}$. Note that this is not equivalent to optimizing Eq. (3) as the max operator is over paths whose expected energy gain is computed given H_t .

Results. We perform 10 simulations of the wind field, each containing 20 gliders, making 200 gliders in total. We simulate for a duration of 500 s. We first consider gliders navigating in wind fields that remain constant in time, which have spatial correlations specified by $K(r, 0)$ in Eq. (2). In Fig. 2, we show a sample path of the glider along with the trajectories explored at a single instance of MCTS along its path. The glider exhibits distinct bouts of localized spiraling behavior during which they gain height, similar to soaring patterns observed during thermal soaring. These bouts of spiraling punctuated by flat traversals can be intuitively viewed as ‘foraging’ behavior: the ascending currents that drove spiraling behavior expire when the height gained in these bouts exceeds the typical correlation length ($\simeq 50$ m). Subsequently, the glider traverses less valuable regions of the flow towards new energy-rich regions.

Over time scales much longer than than 1 s, since v remains comparable to v_c , energy extracted from the flow is primarily converted to potential energy (see SM.VII). We thus use the climb rate c (Fig. 3) of the gliders to measure the efficiency of energy extraction. We simulate gliders that execute random actions and gliders that have full information of the flow to obtain lower (c_{\min}) and upper bounds (c_{\max}) respectively of the climb rate. The efficiency of energy extraction is defined as $\eta(c) = (c - c_{\min}) / (c_{\max} - c_{\min})$. The efficiency estimates the value provided by global knowledge of the flow versus the local information acquired along the glider’s path. Gliders with full information of this frozen flow show positive climb rates $c_{\max} = 0.5$ m/s, implying that energy-positive flight is feasible if gliders have sufficient information and the flow has sufficiently long correlations. Gliders that use partial information—i.e., the measured wind velocities along their past trajectory—also show positive climb rates (Fig. 3, SM.VII) with $\eta = 0.77, 0.81$ for memory sizes $M = 20, 50$ respectively. The information is stored in the memory every $\Delta t = 0.5$ s and the memory duration is $t_m = M\Delta t$.

Next, we consider a dynamic wind field that has temporal correlations denoted by τ_c . The frozen field considered above is obtained from the limit $\tau_c \rightarrow \infty$. For smaller τ_c , the reduced temporal correlations imply that past experience is less informative when picking out the optimal path via MCTS and thus the average climb rate decreases. Moreover, the intuition behind the spiraling bouts suggests that each bout will last at most $\sim \tau_c$, thereby reducing the time spent in energy-rich regions of the flow. When τ_c is sufficiently small, past information




FIG. 3. Gliders’ altitude gain in a frozen wind field, using the partial information acquired from the measured wind histories, using different memory durations $t_m = [2\text{ s}, 5\text{ s}, 10\text{ s}, 15\text{ s}, 20\text{ s}, 25\text{ s}]$. All the gliders initially descend, but start to gain altitude after building a better model for the wind field. The slope of the altitude gain from 200 s onward is measured as the climb rate. See SM.VII for details about the climb rate measurement.

is not predictive of future flow configurations and the net energy gain should reduce to that of a random glider.

Figure 4 shows that energy-positive flight is feasible for $\tau_c \gtrsim 60\text{ s}$ for partial information gliders. The net energy gain increases with τ_c and has weak dependence on memory size provided that $t_m > 10\text{ s}$. The efficiency of the gliders increases monotonically with τ_c even though the climb rate when gliders have perfect information (c_{\max}) decreases with τ_c , implying that the reduced climb rates in short time scale flows is due to the quicker degradation of acquired information rather than fewer sources of energy.

Discussion. Soaring birds are believed to use relatively stable large-scale flow structures, such as thermals or shear layers, to reliably extract energy through efficient sampling and navigation. Our results show that gliders employing a sampling-based planning algorithm can achieve energy-positive flight by exploiting transient fluctuations in a kinematic model of fluid turbulence. Rather than rely on local mechanical cues or optimized flight patterns, which have been shown to be useful for thermal soaring and dynamic soaring respectively, our algorithm uses an adaptive strategy where the recent history of measured wind velocities is used to predict and estimate the energy gained along future trajectories. Unlike those other modes of soaring, a memory of a few tens of seconds is necessary to achieve energy-positive flight, highlighting the importance of storing information to map out the local physical environment.

Most of the energy gained by gliders in our simulations is through localized updrafts, that is, the second term in the right hand side of Eq. (1) rather than wind gradients. The trajectories of gliders reflect patterns of foraging, where efficient exploration of the flow through relatively

straight paths is followed by energy extraction through localized spirals where significant updrafts are present. This picture suggests that the marginal benefit of a memory of a few tens of seconds is possibly a consequence of the correlation length scale of the flow (here $\sim 100\text{ m}$): a glider that travels at 10 m/s requires at least a memory of 10 seconds to acquire two independent samples of the flow and guide its subsequent decisions.

Neutral conditions in the atmospheric boundary layer show correlation timescales of 50 s to 100 s with typical velocity fluctuations of magnitude 1 m/s [34]. Using simulated gliders that have glide-to-sink ratios (15:1) similar to soaring birds, our results suggest that energy-positive flight is indeed feasible at physically relevant timescales. Modern sailplanes travel at much faster speeds, allowing for more efficient exploration, and reach glide-to-sink ratios greater than 40:1 [35]. Sailplanes employing the algorithm devised here can potentially achieve energy-positive flight under stricter restrictions on atmospheric scales. The algorithm applies flows with arbitrary correlation structure, which include environments where updrafts have distinctive spatial arrangements (e.g., cloud streets in the atmospheric boundary layer [36]). An important caveat is that the predictive model relies on the second-order moments of the flow. An interesting open problem is whether non-Gaussian deviations in atmospheric turbulence [37] significantly affect the algorithm’s performance.

G.R. was partially supported by the NSF-Simons Center for Mathematical & Statistical Analysis of Biology at Harvard (award number #1764269) and the Harvard Quantitative Biology Initiative.

* danyunhe@g.harvard.edu

† chr@math.wisc.edu

- [1] R. Baudinette and K. Schmidt-Nielsen, Energy cost of gliding flight in herring gulls, *Nature* **248**, 83 (1974).
- [2] K. Q. Sakamoto, A. Takahashi, T. Iwata, T. Yamamoto, M. Yamamoto, and P. N. Trathan, Heart rate and estimated energy expenditure of flapping and gliding in black-browed albatrosses, *J. Exp. Biol.* **216**, 3175 (2013).
- [3] V. A. Tucker, Metabolism during flight in the laughing gull, *Larus atricilla*, *Am. J. Physiol.* **222**, 237 (1972).
- [4] P. Henningsson and A. Hedenström, Aerodynamics of gliding flight in common swifts, *J. Exp. Biol.* **214**, 382 (2011).
- [5] C. Harvey and D. J. Inman, Aerodynamic efficiency of gliding birds vs comparable UAVs: a review, *Bioinspir. Biomim.* **16**, 031001 (2021).
- [6] G. K. Taylor, K. V. Reynolds, and A. L. Thomas, Soaring energetics and glide performance in a moving atmosphere, *Philos. T. Roy. Soc. B* **371**, 20150398 (2016).
- [7] C. D. Cone, Thermal soaring of birds, *Am. Sci.* **50**, 180 (1962).
- [8] H. D. Shannon, G. S. Young, M. A. Yates, M. R. Fuller, and W. S. Seegar, American white pelican soaring flight times and altitudes relative to changes in thermal depth

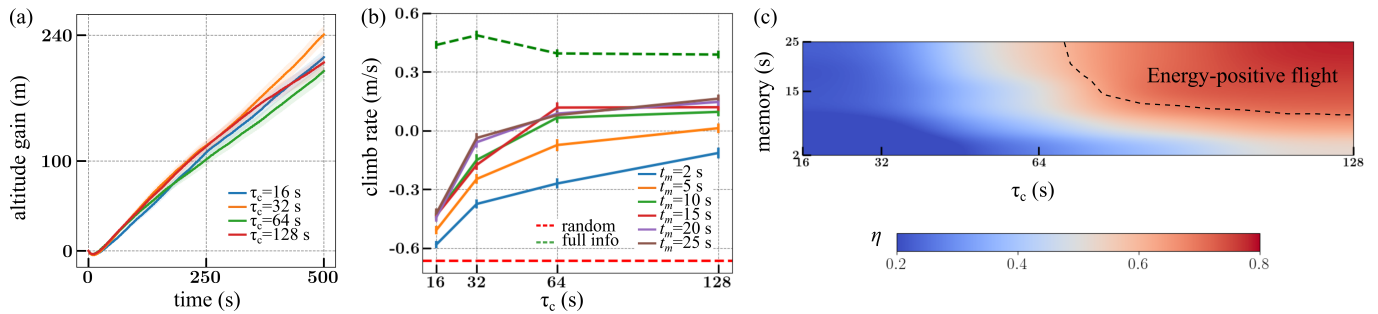


FIG. 4. (a) The average altitude gain of full information gliders in dynamic wind fields with $\tau_c = 16$ s, 32 s, 64 s, 128 s. (b) The summary of climb rates with memory $t_m = 5$ s, 10 s, 15 s, 20 s, 25 s and $\tau_c = 16$ s, 32 s, 64 s, 128 s. The climb rates of full information gliders are shown as the upper bound and the climb rates of random action gliders are shown as the lower bound. (c) Efficiency of energy extraction $\eta(c)$ as a function of t_m and τ_c . The energy-positive flight regime is where $c > 0$.

and intensity, *The Condor* **104**, 679 (2002).

- [9] B. Woodward, A theory of thermal soaring, *OSTIV Publications* **4**, 159 (1956).
- [10] H. J. Williams, O. Duriez, M. D. Holton, G. Dell’Omo, R. P. Wilson, and E. L. Shepard, Vultures respond to challenges of near-ground thermal soaring by varying bank angle, *J. Exp. Biol.* **221**, jeb174995 (2018).
- [11] R. Harel, N. Horvitz, and R. Nathan, Adult vultures outperform juveniles in challenging thermal soaring conditions, *Scientific Reports* **6**, 1 (2016).
- [12] P. L. Richardson, E. D. Wakefield, and R. A. Phillips, Flight speed and performance of the wandering albatross with respect to wind, *Movement Ecology* **6**, 1 (2018).
- [13] G. D. Bousquet, M. S. Triantafyllou, and J.-J. E. Slotine, Optimal dynamic soaring consists of successive shallow arcs, *J. R. Soc. Interface* **14**, 20170496 (2017).
- [14] J. A. Kempton, J. Wynn, S. Bond, J. Evry, A. L. Fayet, N. Gillies, T. Guilford, M. Kavelaars, I. Juarez-Martinez, O. Padget, *et al.*, Optimization of dynamic soaring in a flap-gliding seabird affects its large-scale distribution at sea, *Science Advances* **8**, eabo0200 (2022).
- [15] Y. J. Zhao, Optimal patterns of glider dynamic soaring, *Optimal control applications and methods* **25**, 67 (2004).
- [16] J. M. Mallon, K. L. Bildstein, and T. E. Katzner, In-flight turbulence benefits soaring birds, *The Auk: Ornithological Advances* **133**, 79 (2016).
- [17] C. Patel, H.-T. Lee, and I. Kroo, Extracting energy from atmospheric turbulence with flight tests, *Technical Soaring* **33**, 100 (2009).
- [18] K. V. Reynolds, A. L. Thomas, and G. K. Taylor, Wing tucks are a response to atmospheric turbulence in the soaring flight of the steppe eagle *Aquila nipalensis*, *J. R. Soc. Interface* **11**, 20140645 (2014).
- [19] K. M. Laurent, B. Fogg, T. Ginsburg, C. Halverson, M. J. Lanzone, T. A. Miller, D. W. Winkler, and G. P. Bewley, Turbulence explains the accelerations of an eagle in natural flight, *Proc. Natl. Acad. Sci.* **118**, e2102588118 (2021).
- [20] G. Reddy, A. Celani, T. J. Sejnowski, and M. Vergassola, Learning to soar in turbulent environments, *Proc. Natl. Acad. Sci.* **113**, E4877 (2016).
- [21] G. Reddy, J. Wong-Ng, A. Celani, T. J. Sejnowski, and M. Vergassola, Glider soaring via reinforcement learning in the field, *Nature* **562**, 236 (2018).
- [22] J. Wharington, *Autonomous control of soaring aircraft by reinforcement learning*, Ph.D. thesis, Royal Melbourne Institute of Technology (Australia) (1998).
- [23] M. J. Allen and V. Lin, *Guidance and control of an autonomous soaring UAV*, Tech. Rep. (2007).
- [24] J. J. Chung, N. R. Lawrance, and S. Sukkarieh, Learning to soar: Resource-constrained exploration in reinforcement learning, *Int. J. Robot Res.* **34**, 158 (2015).
- [25] Rayleigh, The soaring of birds, *Nature* **27**, 534 (1883).
- [26] T. D. Woodbury, C. Dunn, and J. Valasek, Autonomous soaring using reinforcement learning for trajectory generation, in *52nd Aerospace Sciences Meeting* (2014) p. 0990.
- [27] Z. Ákos, M. Nagy, S. Leven, and T. Vicsek, Thermal soaring flight of birds and unmanned aerial vehicles, *Bioinspir. Biomim.* **5**, 045003 (2010).
- [28] S. A. Bollt and G. P. Bewley, How to extract energy from turbulence in flight by fast tracking, *J. Fluid Mech.* **921** (2021).
- [29] V. A. Tucker and G. C. Parrott, Aerodynamics of gliding flight in a falcon and other birds, *J. Exp. Biol.* **52**, 345 (1970).
- [30] N. Horvitz, N. Sapir, F. Liechti, R. Avissar, I. Mahrer, and R. Nathan, The gliding speed of migrating birds: slow and safe or fast and risky?, *Ecol. Lett.* **17**, 670 (2014).
- [31] J. C. H. Fung, J. C. R. Hunt, N. A. Malik, and R. Perkins, Kinematic simulation of homogeneous turbulence by unsteady random Fourier modes, *J. Fluid Mech.* **236**, 281 (1992).
- [32] M. Bardi, I. C. Dolcetta, *et al.*, *Optimal control and viscosity solutions of Hamilton–Jacobi–Bellman equations*, Vol. 12 (Springer, 1997).
- [33] C. B. Browne, E. Powley, D. Whitehouse, S. M. Lucas, P. I. Cowling, P. Rohlfshagen, S. Tavener, D. Perez, S. Samothrakis, and S. Colton, A survey of Monte Carlo tree search methods, *IEEE T. Comp. Intel. AI* **4**, 1 (2012).
- [34] D. H. Lenschow and B. B. Stankov, Length scales in the convective boundary layer, *J. Atmos. Sci.* **43**, 1198 (1986).
- [35] F. A. Administration and U. S. F. S. Service, *Glider Flying Handbook* (Aviation Supplies & Academics, 2004).
- [36] F. Irving, Cloud-street flying, *Technical Soaring* **3**, 1 (1973).
- [37] G. G. Katul, M. B. Parlange, and C. R. Chu, Intermittency, local isotropy, and non-gaussian statistics in atmospheric surface layer turbulence, *Phys. Fluids* **6**, 2480 (1994).

Supplemental Material: Energy-positive soaring using transient turbulent fluctuations

Danyun He*

*John A. Paulson School of Engineering and Applied Sciences,
Harvard University, Cambridge, MA 02138, USA*

Gautam Reddy

*Physics & Informatics Laboratories, NTT Research, Inc., Sunnyvale, CA 94085, USA
Center for Brain Science, Harvard University, Cambridge, MA 02138, USA and
NSF-Simons Center for Mathematical & Statistical Analysis of Biology,
Harvard University, Cambridge, MA 02138, USA*

Chris H. Rycroft[†]

*Department of Mathematics, University of Wisconsin–Madison, Madison, WI 53706, USA and
Computational Research Division, Lawrence Berkeley Laboratory, Berkeley, CA 94720, USA*

I. OVERVIEW

This document contains technical details for simulating a glider moving in a model for a turbulent flow to maximize its energy gain. In Sec. II, we introduce the governing equations for a glider moving in a wind field. In Sec. III we describe the model for the turbulent wind field, comprised of stochastically evolving Fourier modes. In Sec. IV we then describe the Gaussian process regression (GPR) that the glider can use to predict the wind field based on its history of observations. In Sec. V we introduce the Monte Carlo tree search (MCTS), which the glider uses to control its flight to maximize energy gain.

Several aspects of the simulation—notably the evolution of the turbulent wind field, and the large number of GPR evaluations required for MCTS—are computationally intensive. Our simulation is therefore written as a custom C++ code, making use of the OpenMP library to accelerate many aspects of the computation using multithreading. Throughout this document, we discuss algorithmic complexity, details of the implementation, and numerical tests that we have performed. In Sec. VI we describe the specific set of simulations that we ran, and in Sec. VII we provide tables of results to supplement the data available in the main paper.

II. GLIDER AERODYNAMICS

A. Governing equations for the glider

We simulate a glider moving in a global (x, y, z) coordinate system where the z axis points upward. The glider moves with time-dependent position $\mathbf{r}(t)$ and velocity $\mathbf{u}(t)$ relative to the ground, experiencing a local wind velocity $\mathbf{w}(t)$. The glider’s motion relative to the wind is therefore given by $\mathbf{v} = \mathbf{u} - \mathbf{w}$, and the airspeed is given by $v = \|\mathbf{v}\|$. As shown in Fig. 1, there are three forces acting on the glider: lift, drag, and weight. Lift is a component of the aerodynamic force, and is directed perpendicular to the flight direction. The magnitude of the lift force is given by

$$L = \frac{1}{2}c_L\rho Sv^2, \quad (1)$$

where ρ is the density of air, S is the surface area of the wing, v is the airspeed and c_L is the dimensionless lift coefficient which has complex dependencies on wing shape. The drag is another component of the aerodynamic force, directed opposite to the flight direction. It scales similarly to lift, and has magnitude

$$D = \frac{1}{2}c_D\rho Sv^2 \quad (2)$$

where c_D is a dimensionless drag coefficient. Figure 1 also defines two orthogonal vectors (e_1, e_2) in the horizontal plane, with e_1 pointing in the same direction in the xy -plane as the glider body. Note however that the glider may also be tilted, so that its direction vector has an additional vertical component.

As shown in Fig. 2, the glider’s orientation can be described using three angles. The bank angle μ sets the tilt of the wing surfaces from vertical, and can be controlled by the glider itself. The glide angle γ measures the angle of the glider’s motion from horizontal and is given by $\gamma = \sin^{-1} v_z/v$. The azimuth angle measures the direction of the glider in the horizontal plane and is determined so that $(v_x, v_y) = (v \cos \gamma \cos \psi, v \cos \gamma \sin \psi)$.

* danyunhe@g.harvard.edu

† chr@math.wisc.edu

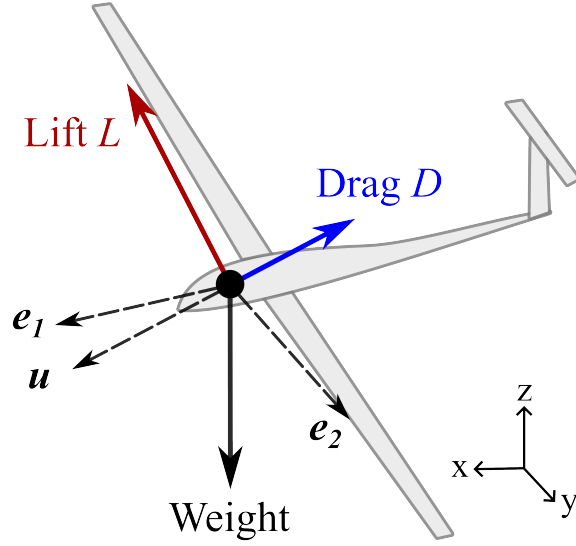


FIG. 1. Diagram of the glider, marking the directions of the three forces, lift L , drag D , and weight that it experiences. The global coordinate system is (x, y, z) , with the positive z direction pointing vertically upwards. The e_1 vector points horizontally in the direction that the glider faces, although the glider direction may also have an additional z component. The vector \mathbf{u} is the glider's velocity relative to the ground. The vector e_2 is horizontal, and is perpendicular to e_1 .




FIG. 2. Diagrams of the glider from several different directions using the global (x, y, z) coordinate system and (e_1, e_2) coordinate system defined in Fig. 1. The angles determine the glider's orientation. (a) Bank angle μ is how much the glider rolls from the vertical. (b) Glide angle γ is the angle between the heading and the horizontal. (c) Azimuth angle ψ is the angle between the projection of the air velocity vector onto the horizontal plane and the x -axis.

With these definitions, the glider's governing equations are given by

$$m\dot{u}_x = L\cos(\mu)\sin(\gamma)\cos(\psi) - L\sin(\mu)\sin(\psi) - D\cos(\gamma)\cos(\psi), \quad (3)$$

$$m\dot{u}_y = L\cos(\mu)\sin(\gamma)\sin(\psi) + L\sin(\mu)\cos(\psi) - D\cos(\gamma)\sin(\psi), \quad (4)$$

$$m\dot{u}_z = L\cos(\mu)\cos(\gamma) + D\sin(\gamma) - mg, \quad (5)$$

$$\dot{\mathbf{r}} = \mathbf{u}, \quad (6)$$

where m is the mass of the glider and g is the gravitational acceleration. The energy is given by

$$E = \frac{1}{2}mv^2 - m\mathbf{g} \cdot \mathbf{z}. \quad (7)$$

Taking the time derivative of Eq. (13) and substituting in Eqs. (9)–(12) shows that the change in energy is given by

$$\begin{aligned} \dot{E} &= m\mathbf{v} \cdot \dot{\mathbf{v}} - m\mathbf{g} \cdot \dot{\mathbf{u}} \\ &= m\mathbf{v} \cdot (\dot{\mathbf{u}} - \dot{\mathbf{w}}) - m\mathbf{g} \cdot (\mathbf{v} + \mathbf{w}) \\ &= \mathbf{v} \cdot (m\dot{\mathbf{u}} - m\mathbf{g}) - m\mathbf{g} \cdot \mathbf{w} - m\mathbf{v} \cdot \dot{\mathbf{w}} \\ &= -vD + mgw_z - m\mathbf{v} \cdot \dot{\mathbf{w}}, \end{aligned} \quad (8)$$

which is given as Eq. (1) in the main text. As described in the main text [1], the first term in this equation corresponds to energy loss by drag. The other two terms in this equation correspond to energy loss by drag, energy gain from altitude increase and energy gain from local wind gusts, both of which are exploited by soaring birds.

B. Non-dimensionalized equations

Define the speed scale $v_c = \sqrt{2mg/\rho S}$ and the corresponding time scale $t_c = v_c/g$ and length scale $l_c = v_c^2/g$. We non-dimensionalize all velocities, lengths and times with respect to these three scales. Retaining the same notation, Eqs. (3)–(6) become

$$\dot{u}_x = c_L v^2 \cos(\mu) \sin(\gamma) \cos(\psi) - c_L v^2 \sin(\mu) \sin(\psi) - c_D v^2 \cos(\gamma) \cos(\psi), \quad (9)$$

$$\dot{u}_y = c_L v^2 \cos(\mu) \sin(\gamma) \sin(\psi) + c_L v^2 \sin(\mu) \cos(\psi) - c_D v^2 \cos(\gamma) \sin(\psi), \quad (10)$$

$$\dot{u}_z = c_L v^2 \cos(\mu) \cos(\gamma) + c_D v^2 \sin(\gamma) - 1, \quad (11)$$

$$\dot{\mathbf{r}} = \mathbf{u}. \quad (12)$$

We use Eqs. (9)–(12) in the numerical implementation. The non-dimensionalized energy is given by dividing Eq. (7) by a factor of mv_c^2 , so that

$$\varepsilon = \frac{1}{2}v^2 + z. \quad (13)$$

The non-dimensionalized rate of energy gain is

$$\dot{\varepsilon} = -c_D v^3 + w_z - \mathbf{v} \cdot \dot{\mathbf{w}}. \quad (14)$$

III. SIMULATIONS OF TURBULENT FLOW

A. Model of turbulent flow using Fourier modes

For the wind field $\mathbf{w}(\mathbf{r}, t)$, we consider a homogeneous and isotropic turbulent flow that can be described by a Fourier transform,

$$\mathbf{w}(\mathbf{r}, t) = (2\pi)^{3/2} \int_{-\infty}^{\infty} d^3\mathbf{k} e^{i\mathbf{k} \cdot \mathbf{r}} \tilde{\mathbf{w}}(\mathbf{k}, t), \quad (15)$$

where $\tilde{\mathbf{w}}(\mathbf{k}, t)$ is the complex-valued amplitude of the wave with wave number \mathbf{k} at time t . We simulate a turbulent field that reproduces the second-order Kolmogorov statistics [2] where the energy contained in

waves of magnitude $k = \|\mathbf{k}\|$ scales like $E(k) \sim k^{-5/3}$. In our numerical implementation, using the non-dimensionalization introduced in Sec. II B, the wind velocity is computed in a periodic box $[0, 2\pi]^3$ using a discrete $N \times N \times N$ complex Fourier modes $\tilde{\mathbf{w}}_{\alpha\beta\zeta}$. Define a lattice spacing $h = 2\pi/N$. For a lattice point $\mathbf{r} = (hj, hk, hl)$, the wind velocity is given by the discrete Fourier transform

$$\mathbf{w}(\mathbf{r}, t) = \sum_{\alpha=0}^{N-1} \sum_{\beta=0}^{N-1} \sum_{\zeta=0}^{N-1} \tilde{\mathbf{w}}_{\alpha\beta\zeta}(t) e^{2\pi i(j\alpha + k\beta + l\zeta)/N}. \quad (16)$$

We treat the complex Fourier modes as indexed periodically, so that $\tilde{\mathbf{w}}_{(\alpha+N),\beta,\zeta} = \tilde{\mathbf{w}}_{\alpha\beta\zeta}$, with similar relations for β and ζ . In order to ensure that $\tilde{\mathbf{w}}(\mathbf{r}, t)$ is real-valued, the discrete Fourier coefficients must satisfy

$$\tilde{\mathbf{w}}_{\alpha\beta\zeta} = \tilde{\mathbf{w}}_{-\alpha,-\beta,-\zeta}^*. \quad (17)$$

Due to the periodicity of the complex Fourier modes, the sums in Eq. (16) can be reordered. Define $M = \lfloor \frac{N}{2} \rfloor$, and

$$f(n) = \begin{cases} 1 & \text{if } |n| = N/2, \\ 0 & \text{if } |n| < N/2. \end{cases} \quad (18)$$

Then for a lattice point $\mathbf{r} = (hj, hk, hl)$, Eq. (16) is equivalent to

$$\mathbf{w}(\mathbf{r}, t) = \sum_{\alpha=-M}^M \sum_{\beta=-M}^M \sum_{\gamma=-M}^M f(\alpha)f(\beta)f(\gamma) \tilde{\mathbf{w}}_{\alpha\beta\zeta}(t) e^{2\pi i(j\alpha + k\beta + l\zeta)/N}. \quad (19)$$

When N is odd, each sum has exactly N terms. When N is even, each sum has exactly $N + 1$ terms, but f function causes the two extremal indices at $\pm M$ to be counted with half weighting. Those two extremal indices have the same Fourier mode coefficient, due to the periodicity. Thus in both cases the sums evaluate the same terms as in Eq. (16).

Equations (16) & (19) are equivalent at lattice points, but this relies on cancellation of factors of $e^{2\pi i}$. At off-lattice points they are not the same, since Eq. (16) involves a sum over rapidly oscillating exponentials. By contrast, Eq. (19) provides a representation of the Fourier exponentials in the lowest frequencies. Since the gliders move across arbitrary locations, we therefore make use of Eq. (19) to evaluate the wind field at any position \mathbf{r} .

We now identify $\tilde{\mathbf{w}}(\mathbf{k}) = \tilde{\mathbf{w}}_{\alpha\beta\zeta}$ for $\mathbf{k} = (\alpha, \beta, \zeta)$. Following the model of Fung et al. [2] we assume that each Fourier mode evolves as an Ornstein–Uhlenbeck process with time scale $\tau(k)$ so that

$$d\tilde{\mathbf{w}}(\mathbf{k}, t) = -\tilde{\mathbf{w}}(\mathbf{k}, t) \frac{dt}{\tau(k)} + a(k) dW(t), \quad (20)$$

where $W(t)$ denotes the Wiener process. To match Kolmogorov scaling, we require $\tau(k) \sim k^{-2/3}$. $a(k)$ is set such that the energy spectrum $E(k) \sim k^{-5/3}$. The specific expressions for $a(k)$ and $\tau(k)$ are derived below in Section III C.

B. Numerical implementation

To visualize the wind field, it is useful to perform a standard discrete Fourier transform at lattice locations using Eq. (16). This was used to create the wind plots in Fig. 1 in the main text. To perform this, we make use of the FFTW library, and in particular the `c2r` transforms for transforming complex Fourier modes into a real field. Because of the restriction on the modes given by Eq. (17), FFTW represents the modes

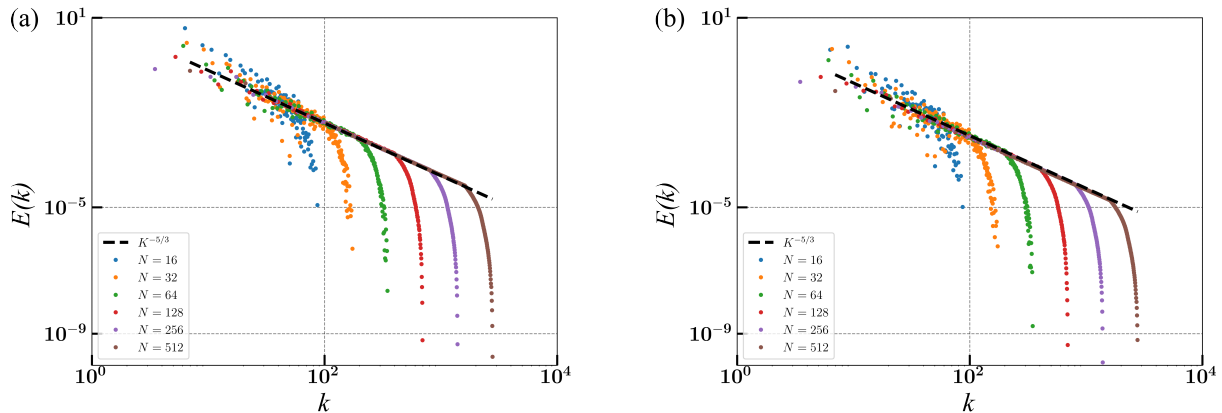


FIG. 3. The energy spectrum in simulations using different number of modes ($N = [16, 32, 64, 128, 256, 512]$) captures Kolmogorov’s 5/3 law. (a) The initial wind field at $t = 0$. (b) The wind field at $t = 100$ s.

in a $N \times N \times \lfloor \frac{N+1}{2} \rfloor$ array, where the other modes are implicitly defined. FFTW represents each complex number using a custom `fftw_complex` data type, consisting of an array of two double precision floating point numbers (occupying 16 bytes). We therefore use this three-dimensional array as primary storage for the Fourier modes that describe the wind field.

We use a custom `fftw_plan` to perform the fast Fourier transforms for all three velocity components together, which due to vectorization provides some improvements in speed over doing three sequential computations. FFTW also performs multithreading computation and evaluates the complete three dimensional discrete Fourier transform in $O(N^3 \log N)$ time.

To evaluate the wind field at a single off-lattice glider position, we wrote a custom routine for computing Eq. (19), which can be performed in $O(N^3)$ time. Because all the modes make independent computations to the sum in Eq. (19), this computation can be efficiently multithreaded with OpenMP. Since this only has to compute a single velocity instead of a complete three-dimensional grid of velocities, it is considerably faster than the FFTW computation. In some cases, we consider a set of G gliders within the same wind field, and we wrote an additional routine for computing G velocities simultaneously. Since all of the $O(N^3)$ Fourier modes only need to be read from memory once, this routine is faster than performing G individual wind evaluations.

Finally, we wrote a routine for updating the Fourier modes according to Eq. (20). We make use of the GNU Scientific Library (GSL) for computing the Gaussian random variables, with the Tausworthe random number generator [3, 4]. This update can be multithreaded efficiently, with each thread using its own generator initialized with different seeds.

In Fig. 3, we show that the energy spectrum of the wind fields at the initial state and a later state $t = 100$ s. Since the wind field is initialized at steady state, it satisfies Kolmogorov’s power law that $E(k) \sim k^{-5/3}$. The later wind field also satisfies Kolmogorov’s law, which demonstrates that that our stochastic differential equation in Eq. (20) preserves the modes at equilibrium. As the number of modes N is increased, Kolmogorov’s law is extended to a larger range of k .

C. Spatio-temporal covariance

In the later sections, we are going to infer the wind information at unknown locations and times with the partial knowledge of the field. Thus, we need to know the correlation of the wind function. To compute the

spatio-temporal covariance function, we first write

$$\begin{aligned}\langle w(0,0), w(r,t) \rangle &= (2\pi)^{-3} \int \int d^3\mathbf{k} d^3\mathbf{k}' e^{i\mathbf{k}'\cdot\mathbf{r}} \langle \tilde{w}(\mathbf{k},0), \tilde{w}(\mathbf{k}',t) \rangle \\ &= (2\pi)^{-3} \int \int d^3\mathbf{k} d^3\mathbf{k}' e^{i\mathbf{k}'\cdot\mathbf{r}} (12\delta(\mathbf{k}+\mathbf{k}')e(k)e^{-|t|/\tau(k)}\end{aligned}\quad (21)$$

where the second step comes from

$$\begin{aligned}\langle \tilde{w}(\mathbf{k},0), \tilde{w}(\mathbf{k}',t) \rangle &= \langle \tilde{w}_R(\mathbf{k},0), \tilde{w}_R(\mathbf{k}',t) \rangle - \langle \tilde{w}_I(\mathbf{k},0), \tilde{w}_I(\mathbf{k}',t) \rangle \\ &\quad + i\langle \tilde{w}_R(\mathbf{k},0), \tilde{w}_I(\mathbf{k}',t) \rangle + i\langle \tilde{w}_I(\mathbf{k},0), \tilde{w}_R(\mathbf{k}',t) \rangle \\ &= (\delta(\mathbf{k}-\mathbf{k}') + \delta(\mathbf{k}+\mathbf{k}'))6e(k)e^{-|t|/\tau(k)} \\ &\quad - (\delta(\mathbf{k}-\mathbf{k}') - \delta(\mathbf{k}+\mathbf{k}'))6e(k)e^{-|t|/\tau(k)} + 0 + 0 \\ &= 12\delta(\mathbf{k}+\mathbf{k}')e(k)e^{-|t|/\tau(k)}\end{aligned}\quad (22)$$

Integrating out \mathbf{k}' gives

$$\begin{aligned}\langle w(0,0), w(r,t) \rangle &= \frac{3}{2\pi^3} \int d^3\mathbf{k} e(k) e^{-i\mathbf{k}\cdot\mathbf{r}} e^{-|t|/\tau(k)} \\ &= \frac{3}{\pi^2} \int_{-1}^1 d(\cos\theta) \int_{k_{\min}}^{k_{\max}} dk k^2 e(k) e^{ikr\cos\theta} e^{-|t|/\tau(k)}.\end{aligned}\quad (23)$$

The largest length-scale L is set to unity, then we have the smallest wave-number $k_{\min} = \sqrt{3}\pi/L = \sqrt{3}\pi$. The largest wave-number is determined by the smallest resolved length-scale, which is $1/N$ in an N^3 grid, therefore, we have the largest wave-number $k_{\max} = \sqrt{3}\pi/(1/N) = \sqrt{3}\pi N$. Let $k^2 e(k) = \alpha k^{-5/3}$ where α will be set such that the variance of each component is 1. Then we have

$$\begin{aligned}\langle w(0,0), w(0,0) \rangle = 1 &= \frac{3}{\pi^2} \int_{-1}^1 d(\cos\theta) \int_{k_{\min}}^{k_{\max}} dk \alpha k^{-5/3} \\ &= \frac{6\alpha}{\pi^2} \int_{k_{\min}}^{k_{\max}} dk k^{-5/3} \\ &= \frac{9\alpha}{\pi^2} (k_{\min}^{-2/3} - k_{\max}^{-2/3})\end{aligned}\quad (24)$$

which gives

$$\begin{aligned}\alpha &= \frac{\pi^2}{9} (k_{\min}^{-2/3} - k_{\max}^{-2/3})^{-1}, \\ e(k) = \alpha k^{-11/3} &= \frac{\pi^2}{9} k^{-11/3} (k_{\min}^{-2/3} - k_{\max}^{-2/3})^{-1}, \\ a(k) = \sqrt{4e(k)/\tau(k)} &= \frac{2}{3} 3^{-1/6} \pi^{1/3} k^{-3/2} \sqrt{(k_{\min}^{-2/3} - k_{\max}^{-2/3})^{-1}}.\end{aligned}\quad (25)$$

The covariance function for each component is $K(r,t) = \langle w(0,0), w(\mathbf{r},t) \rangle$, and

$$K(r,t) = \frac{6\alpha}{\pi^2} \int_{k_{\min}}^{k_{\max}} dk k^{-5/3} \text{sinc}(kr) e^{-|t|/\tau(k)}.\quad (26)$$

Due to the translation invariance of the wind field, for any two locations r_1, r_2 , times t_1, t_2 , we compute $r = \|\mathbf{r}_1 - \mathbf{r}_2\|, t = |t_1 - t_2|$ and apply Eq. (26) to get their correlation.

The analytical correlation function $K(r,t)$ is illustrated in Fig. 4, where the correlation decreases exponentially as difference in time or distance increases. There is significant correlation when $r < 0.5$ and $t < 1$. The correlation of the wind field simulation $\hat{K}(r,t)$ matches the analytical values $K(r,t)$.




FIG. 4. Illustration of the kernel function. (a) The correlation function with respect to time t and distance r when $\tau_c = 1$ s. (b) The correlation function with respect to r at $t = 0$ and $t = 1$ when $\tau_c = 1$ s. The correlation of the simulated wind field $\hat{K}(r,t)$, which is computed using 128 wind fields, each has 1024 samples at random locations. (c) The correlation function with respect to r at $t = 10$ s when $\tau_c = [1 \text{ s}, 16 \text{ s}, 32 \text{ s}, 64 \text{ s}, 128 \text{ s}]$.

IV. GAUSSIAN PROCESS REGRESSION

A. Wind prediction with GPR

We define the mean function $m(\mathbf{x})$ and the covariance function $k(\mathbf{x}, \mathbf{x}')$ of a Gaussian process $f(\mathbf{x})$ as [5]

$$m(\mathbf{x}) = E[f(\mathbf{x})], \quad (27)$$

$$f(\mathbf{x}) \sim GP(m(\mathbf{x}), k(\mathbf{x}, \mathbf{x}')). \quad (28)$$

Suppose we have an observed dataset X of size M , and define $f = f(X)$. Let f^* be unobserved data at X^* . Then the joint distribution of f and f^* is

$$\begin{bmatrix} f \\ f^* \end{bmatrix} \sim N \left(\begin{bmatrix} m(X) \\ m(X^*) \end{bmatrix}, \begin{bmatrix} K(X, X) & K(X, X^*) \\ K(X^*, X) & K(X^*, X^*) \end{bmatrix} \right) \quad (29)$$

Conditioning the joint Gaussian prior distribution on the observations, we have

$$f^* | f, X, X^* \sim N(K(X^*, X)K(X, X)^{-1}f, K(X^*, X^*) - K(X^*, X)K(X, X)^{-1}K(X, X^*)). \quad (30)$$

Since our wind components are generated from many random modes that evolve as a Gaussian process, we use Gaussian process regression to perform wind inference, that is, we can apply Eq. (30) to perform wind predictions using a history of observations. We calculate the kernel function directly from the wind model Eq. (26). Figure 5 gives an example of the wind predictions with sparse sampling in the wind field. It recovers the large scale features in the wind field. To validate the accuracy of the GPR prediction, we set 10 frozen wind fields, and let 20 gliders randomly explore for $t_m = 25$ s, with the glider storing information $(\mathbf{w}, \mathbf{r}, t)$ in memory every $\Delta t = 0.5$ s. When the memory is full, we perform predictions after 0.1, 1, and 10 s, where 10 s is the planning horizon of gliders. Figure 6 demonstrates that the GPR prediction is accurate over short distances and times, and has large errors over long times where the correlation is weak. Increasing the duration of memory collection helps improve the performance of the GPR prediction. Increasing the number of modes brings more fine details into the wind field, which makes the prediction harder.

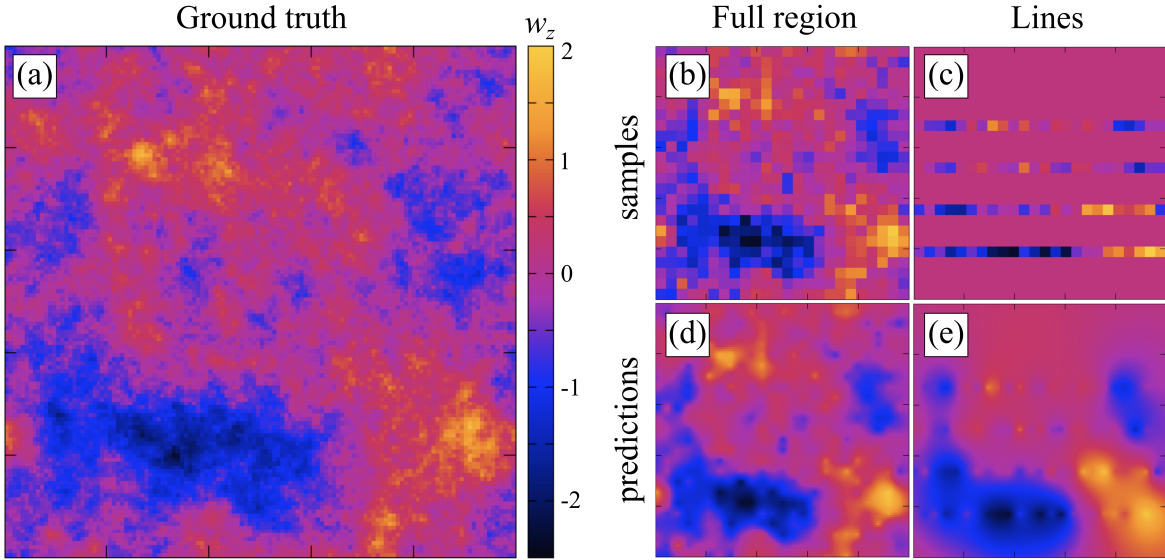


FIG. 5. Illustration of GPR prediction. (a) Vertical wind velocity in 128×128 xz -plane computed using the FFTW library as described in Sec. III B. (b) Coarse 24×24 samples of the wind field (c) Four lines of samples of the wind field on a 4×24 grid. (d) The corresponding GPR predictions of w_z using the samples in (b). (e) The corresponding GPR predictions using the samples in (c).

B. Efficient Gaussian process regression

Write $X = [x_1, x_2, \dots, x_M]$. Evaluating the mean in Eq. (30) requires computing $K(X, X)^{-1}$, which will change on each step. Directly inverting $K(X, X)$ each time, such as by using the LU decomposition [?], will require $O(M^3)$ operations, which will become prohibitively expensive. We therefore accelerate the computation of these inverses.

As gliders explore the field, the new memory replaces the old memory. When a new observation x_{M+1} is added to X , the oldest observation x_1 is deleted, then we have $X' = [x_{M+1}, x_2, \dots, x_M]$. We can compute the new inverse of covariance function $K(X', X')^{-1}$ using $K(X, X)^{-1}$ efficiently by applying the Woodbury formula. Let $A = K(X, X)$, $v = K(x_{M+1}, X)$, and suppose A^{-1} is known. Write

$$K(X', X')^{-1} = (A + ev^T + ve^T)^{-1}.$$

Where e is a unit vector that $e_1 = 1$. According to the Woodbury formula

$$(A + ev^T + ve^T)^{-1} = A^{-1} - A^{-1}U(C + VA^{-1}U)VA^{-1}$$

where

$$U = [e|v], \quad V = U^T, \quad C = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$$

The updating step requires $O(M^2)$ time, which is more efficient than computing the inverse of the matrix directly.

V. MONTE CARLO TREE SEARCH

Monte Carlo tree search is an efficient method for planning into the future and making decisions. The key idea is to randomly build a search tree of possible actions into the future. In the method, many random




FIG. 6. The relative errors of wind prediction on the frozen wind field in the future time. There are 200 gliders randomly travel in the frozen wind field. In the first 25 s, they collect information and store in the memory with $M = 50$. Then they use GPR to do predictions at their current locations after 0.1 s, 1 s and 5 s. The figure shows the relative errors between the wind prediction $\hat{\mathbf{w}}$ and actual wind velocities \mathbf{w} .

play-outs of the search tree are considered, where more promising branches are visited more often, but there is still exploration of less fruitful options. After each play-out, the final reward is used to update the rewards at each node along the path of that play-out. After playing out many simulations, the tree provides an estimation of the best trajectory.

Define the sequence of k actions $A_k = a_1 a_2 a_3 \cdots a_k$, where each action is a change in bank angle. Specifically each a_j will be $\Delta\mu = \pm 10^\circ$, and μ ranges from -40° to 40° . The value function $V(A_k)$ denotes the estimated average energy gain from A_k till terminate. $N(A_k, a_{k+1})$ denotes the number of visited times of a_{k+1} from A_k . For node i of the tree, it stores the information of the action sequence A_i , the value function $V(A_i)$ and the total number of visits $N(A_{i-1}, a_i)$. The UCT algorithm selects the action that maximizes the upper confidence bound [6]

$$U(A_k, a_{k+1}) = V(A_k \cup a_{k+1}) + 2C \sqrt{\frac{2 \ln(\sum_a N(A_k, a))}{N(A_k, a_{k+1})}}. \quad (31)$$

The first term in the UCT bound encourages choosing high-value action, while the second term encourages exploring the less visited actions. $C = 0.5$ is the exploration parameter that can be varied if more exploration is desired.

In one Monte Carlo simulation, a sequence of actions is selected using the UCT algorithm, then the energy gain is estimated using Eq. (14) and propagated back to update $V(A_k)$. To estimate the energy gain, the wind velocities in future times and locations are needed. We perform the GPR prediction to estimate wind and then compute the estimated energy gain. A sufficient amount of Monte Carlo simulations are played out and the action that has the highest value is chosen to be the actual action.

VI. SIMULATION SETUP

We simulate gliders soaring in the wind field. They collect information during exploration, and then use MCTS planning scheme with GPR prediction to choose the optimal action and move forward. The

gliders may know the full region wind field or only the part they have experienced. We define them as a *full information glider* and *partial information glider* accordingly.

A. Glider setup

a. Full information glider. We define the full information glider to be the glider that knows the wind velocities in the whole region at the current time. In the MCTS planning part, it uses the precise wind velocities $\mathbf{w}(\mathbf{r}, t)$ in the current time. When planning, it uses $\mathbf{w}(\mathbf{r}, t)$ to predict $\mathbf{w}(\mathbf{r}, t + \Delta t)$ with Eq. (32), which only involves temporal correlations in the wind field, so that

$$K(0, t) = \frac{|t|}{\tau(k_{\min}) - \tau(k_{\max})} \left(\Gamma \left(-1, \frac{|t|}{\tau(k_{\min})} \right) - \Gamma \left(-1, \frac{|t|}{\tau(k_{\max})} \right) \right). \quad (32)$$

b. Partial information glider. The partial information glider refers to the glider that only knows the wind velocities at its visited locations. It stores limited wind information with memory size $M = 10, 20, 30, 40, 50$. In the planning part, it uses GPR to predict the wind at unknown locations and times using the prior knowledge in memory.

B. Wind setup

We have two wind settings: the frozen wind field and the dynamic wind field. We simulate the wind using $N = 64$ and the box size is $500\text{m} \times 500\text{m} \times 500\text{m}$. In the frozen wind field, we assume the wind does not change over time, which is equivalent to the limit $\tau(k) \rightarrow \infty$. In the dynamic fields, we consider time scales $\tau(k) = 2^i(k/\sqrt{3}\pi)^{-2/3}$, where i are integers from 1 to 8.

C. MCTS setup

We can customize MCTS by choosing the depth of the planning tree and number of simulated paths. The planning horizon is chosen to let the glider sufficiently explore the field. We set the planning horizon to be 10 s for two reasons: first, the planning horizon is sufficient for the glider to make a full circle to explore the field; second, there is reasonable correlation of wind within 10 s (Fig. 4). As shown in Fig. 6, the gliders cannot gather accurate information in after 5 s. We set the number of paths to be 10^4 , which is sufficient to select the optimal action and is computationally efficient.

VII. RESULTS

A. Climb rate and efficiency

To evaluate the performance, we measure the average climb rates of the gliders based on the altitude gained from 200 s to 500 s. We consider the frozen wind fields and dynamic wind fields where $\tau_c = 16\text{s}, 32\text{s}, 64\text{s}, 128\text{s}$. We then have full information gliders and partial information gliders with different memory sizes $M = 4, 10, 20, 30, 40, 50$. Table I summarizes the details of the climb rates.

The full information gliders set an upper bound on the gliders' performance and the sink rate sets the lower bound. Thus, we define the efficiency as

$$\eta(c) = \frac{c - c_{\min}}{c_{\max} - c_{\min}}. \quad (33)$$

Table II presents the efficiencies of partial gliders in dynamic wind fields and frozen wind fields.

TABLE I. Summary of climb rates (m/s) of partial information gliders and full information gliders in dynamic wind fields and frozen wind fields.

$\tau_c(s) \backslash t_m(s)$	random	2	5	10	15	20	25	Full info.
16	-0.67	-0.581	-0.507	-0.428	-0.432	-0.441	-0.422	0.437
32	-0.67	-0.375	-0.247	-0.150	-0.175	-0.060	-0.038	0.486
64	-0.67	-0.270	-0.074	0.065	0.117	0.086	0.079	0.394
128	-0.67	-0.115	0.012	0.095	0.119	0.146	0.162	0.388
frozen	-0.67	-0.024	0.092	0.145	0.161	0.170	0.176	0.398

TABLE II. Summary of efficiencies of partial information gliders in dynamic wind fields and frozen wind fields.

$\tau_c(s) \backslash t_m(s)$	2	5	10	15	20	25
16	7.74%	14.45%	21.66%	21.25%	20.44%	22.20%
32	25.30%	36.37%	44.79%	42.62%	52.65%	54.54%
64	37.35%	55.28%	68.27%	73.12%	70.24%	69.54%
128	52.32%	64.33%	72.24%	74.49%	77.06%	78.62%
frozen	60.38%	71.25%	76.23%	77.81%	78.66%	79.16%

B. Components of energy gain

To investigate the composition of energy gain, we compute each component in Eq. (14). Figure 7 shows that the contributions from wind fluctuations are insignificant. The energy gained is mainly due to upcurrents.

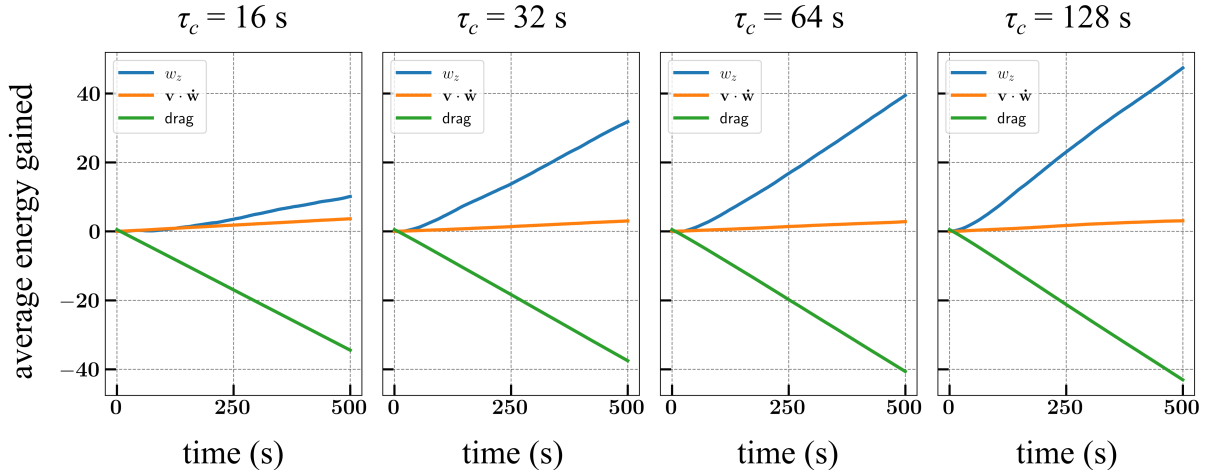


FIG. 7. The components of average energy gained. We simulated 200 gliders for $\tau_c = 16, 32, 64, 128$ s with $M = 40$ and show the average energy gained in Eq. (14) separately. The w_z term dominates the energy contribution.

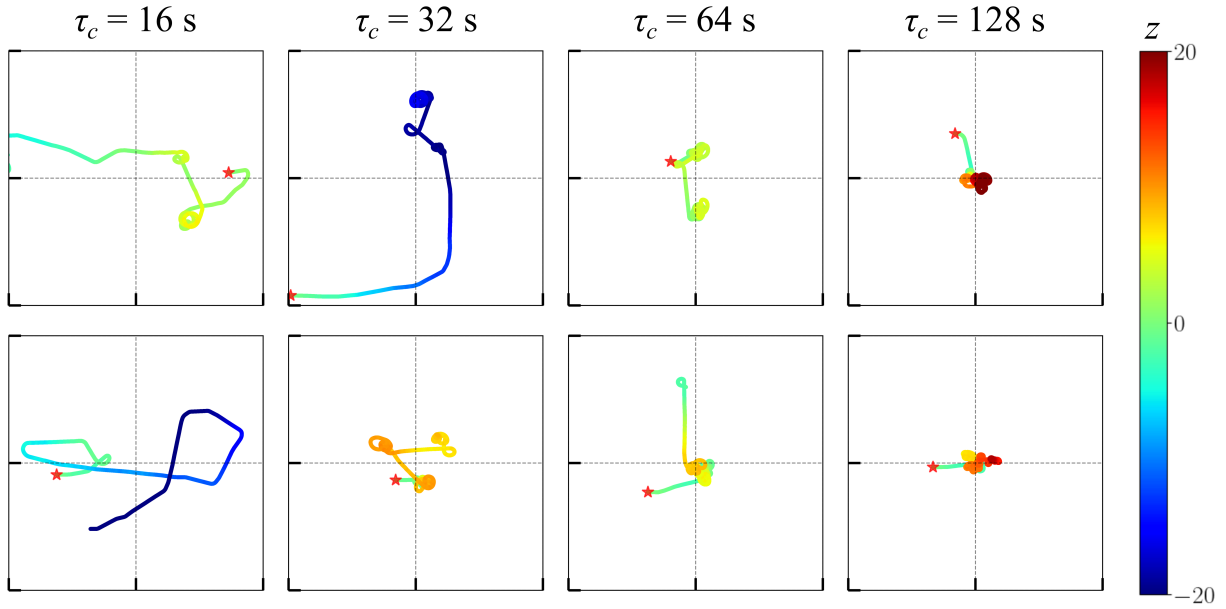


FIG. 8. Examples of gliders trajectories in dynamic wind fields. Top view of gliders $100 \times 100m^2$ xy -plane. The red star indicates the initial positions of the gliders.

C. Gliders' trajectories

The partial information gliders demonstrate different exploration strategies under different field dynamics. As shown in Fig. 8, for more fluctuating fields, the gliders discover and stop at different locations, while in more steady fields, the gliders behave more stably.

TABLE III. Parameters used throughout the paper.

Symbol	Description	Value
Glider parameters		
t	Time	
\mathbf{r}	Glider's position	
\mathbf{u}	Glider's ground velocity	$\ \mathbf{u}_0\ = 10 \text{ m/s}$
\mathbf{w}	Wind velocity	
w	Magnitude of wind velocity, $w = \ \mathbf{w}\ $	$\bar{w} = 0.5 \text{ m/s}$
\mathbf{v}	Glider's air velocity, $\mathbf{v} = \mathbf{u} - \mathbf{w}$	
v	Glider's air speed, $v = \ \mathbf{v}\ $	
ρ	Air density	
S	Surface area of the wing	
c_L	Dimensionless lift coefficient	15
c_D	Dimensionless drag coefficient	1
L	Magnitude of the lift	
D	Magnitude of the drag	
m	Mass of glider	
g	Gravitational acceleration	
μ	Glider's bank angle	$[-40^\circ, 40^\circ], \Delta\mu = \pm 10^\circ$
γ	Angle of the glider's motion from horizontal	
ψ	Glider's azimuth angle	
v_c	Speed scale	$\sqrt{2mg/\rho S}$
t_c	Time scale	v_c/g
l_c	Length scale	v_c^2/g
E	Total energy of glider	
ε	Dimensionless total energy of glider	
G	Number of gliders in the same wind field	20
Wind parameters		
\mathbf{k}	Wave number	
k	Magnitude of wave, $k = \ \mathbf{k}\ $	
k_{\min}	Smallest wave number	$\sqrt{3}\pi$
k_{\max}	Largest wave number	$\sqrt{3}\pi N$
$\tilde{\mathbf{w}}$	Complex-valued amplitude of the wave	
$E(k)$	Energy contained in waves	
N	Number of complex Fourier modes in one dimension	
τ_c	Temporal correlations of the dynamic wind field	16 s, 32 s, 64 s, 128 s
$\tau(k)$	Time scale of the of the stochastic process	$\tau_c(k/\sqrt{3}\pi)^{-2/3}$
$K(r, t)$	Correlation function of each wind component	Eq. (26)
GPR parameters		
M	Memory size	4, 10, 20, 30, 40, 50
Δt	Time interval for one memory storage	0.5 s
t_m	Memory duration	2 s, 5 s, 10 s, 15 s, 20 s, 25 s
MCTS parameters		
H_t	Glider's history at time t	
$V(H_t)$	Maximum expected energy gained given H_t	
d	Planning depth	20
Result parameters		
c	Climb rate of partial information glider from 200 s to 500 s	Table I
c_{\min}	Sink rate	-1/15
c_{\max}	Climb rate of full information glider	Table I
η	Efficiency of partial information glider	Table II

-
- [1] G. K. Taylor, K. V. Reynolds, and A. L. Thomas, Soaring energetics and glide performance in a moving atmosphere, *Philos. T. Roy. Soc. B* **371**, 20150398 (2016).
 - [2] J. C. H. Fung, J. C. R. Hunt, N. A. Malik, and R. Perkins, Kinematic simulation of homogeneous turbulence by unsteady random Fourier modes, *J. Fluid Mech.* **236**, 281 (1992).
 - [3] P. L'Ecuyer, Maximally equidistributed combined Tausworthe generators, *Math. Comp.* **65**, 203 (1996).
 - [4] P. L'Ecuyer, Tables of maximally equidistributed combined LFSR generators, *Math. Comp.* **68**, 261 (1999).
 - [5] Z. Y. Wan and T. P. Sapsis, Reduced-space gaussian process regression for data-driven probabilistic forecast of chaotic dynamical systems, *Physica D* **345**, 40 (2017).
 - [6] L. Kocsis and C. Szepesvári, Bandit based Monte-Carlo planning, in *Machine Learning: ECML 2006: 17th European Conference on Machine Learning Berlin, Germany, September 18-22, 2006 Proceedings 17* (Springer, 2006) pp. 282–293.