

Information Bound on Navigation Speed in Smart Active Matter

Kristian Stølevik Olsen^{1,*}, Mitsusuke Tarama², and Hartmut Löwen¹

¹*Institute for Theoretical Physics II-Soft Matter, Heinrich Heine University Düsseldorf, D-40225 Düsseldorf, Germany*

²*Department of Physics, Kyushu University, Fukuoka 819-0395, Japan*



(Received 3 March 2026; accepted 15 September 2026; published 9 October 2026)

Intelligent behavior in lifelike systems often arises from the ability to gather, process, and act on information. While active matter provides a framework for studying lifelike dynamics, it typically omits internal information processing and decision making. Here we introduce an adaptive active particle model that uses minimal information processing capabilities in order to navigate toward a distant target. By combining renewal-based intermittent motion with the Cramér-Rao inequality, we derive a bound on the navigation speed valid for a wide range of information processing strategies. The framework captures hallmark features of cognitive systems, including optimal sensing durations and a trade-off that balances noise and reliability. Allowing stored information to degrade before action reveals that, although deterioration slows navigation, the trade-off remains governed primarily by external orientational noise and is remarkably insensitive to memory decay.

DOI: [10.1103/PhysRevLett.137.157101](https://doi.org/10.1103/PhysRevLett.137.157101)

Introduction—Living and lifelike systems are remarkable in their ability to perform purposeful behavior in complex or noisy environments. Such behavior emerges from the ability to sense, process, and act upon information about the environment, a crucial property not shared by passive counterparts [1]. Understanding and characterizing how such behavior emerges from the interplay between information and dynamics is central to ongoing efforts to extend the field of active matter into the realm of intelligent systems [2–4].

Active matter provides a powerful framework within statistical physics for modeling the dynamics of self-driven biological agents and is based on the ability to convert energy input into sustained motion [5]. While such energy processing is a crucial property of lifelike systems, so is the ability to process information [6]. Whether in molecular motors adjusting their steps to environmental cues [7,8] or in large animals foraging across complex landscapes [9,10], information underpins the physical behavior and efficiency of these systems.

However, only recently has the field begun to grapple with the informational aspects of behavior. Early developments such as chemotaxis-inspired models laid the groundwork for incorporating goal-oriented navigation into particle dynamics [11–13]. Recent approaches have

explored extensions of active matter by endowing particles with memory, internal states, or adaptive rules, leading to models often referred to as smart or intelligent active matter [14–24]. Machine and reinforcement learning methods have also been used to reveal optimal strategies for navigating complex environments in scenarios where analytical treatment is intractable [25–36]. These efforts aim to bridge the gap between passive response and active decision making. Beyond theoretical insights, they are also relevant to real-world applications such as in autonomous robotics, where processing partial or noisy information in real time to navigate toward a goal is key, as well as to biological systems like insects that rely on incomplete internal representations to adapt their dynamics [37–40].

In parallel to these developments, several disciplines have developed formal frameworks to understand decision making under uncertainty. In quantitative neuroscience and behavioral science, minimal models capture generic features such as the tendency of agents to balance speed and accuracy when making time-sensitive decisions [41–47]. Furthermore, search and navigation strategies based on Bayesian inference have been used to model chemotaxis and other sensory-guided behaviors [48–50]. Yet despite their relevance, such concepts are rarely integrated directly into active matter models in an analytically tractable way, and consequently information is often implicit or absent altogether.

In this Letter, we propose a model that combines classical active matter models and information theory by allowing particles to perform simple inference of a target direction. Within this framework, we demonstrate the existence of an information-constrained bound on the effective navigation speed, originating from the

*Contact author: kristian.olsen@hhu.de

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

Cramér-Rao inequality [51,52], and use this to analyze different information-processing strategies. Generic behavioral features such as the speed-accuracy trade-off [53], reflecting the tension between acting quickly and acting reliably, and the existence of optimal information processing times emerge from the model, demonstrating how minimal inference properties combined with active dynamics can account for generic behavioral features. Despite the simplicity of our model, it reveals rich behaviors and offers a tractable setting for studying how information shapes active dynamics.

Information-based navigation strategies and Cramér-Rao speed limits—We consider a renewal-based navigation strategy where the searcher has no long-term memory of the target direction. Rather, it collects information for a duration τ , drawn from a normalized *renewal time* distribution $\psi(\tau)$, then proceeds by acting on this information and effectively erasing it. This scheme endows the system with a renewal structure that, while sufficient to incorporate and produce realistic effects, also allows analytical analysis. Dynamically, the strategy is biphasic and consists of two repeated phases: (1) runs with information acquisition and (2) informed reorientations (see Fig. 1): 1. *Runs and information acquisition*: During the n th run, which lasts for a random time τ_n drawn from the renewal time density $\psi(\tau_n)$, the particle follows the dynamics of an active Brownian particle in two dimensions [54,55]:

$$\dot{\vec{x}}(t) = v_0 \hat{n}(t), \quad (1)$$

$$\dot{\phi}(t) = \sqrt{2D_r} \xi(t), \quad (2)$$

where v_0 is the self-propulsion speed, and $\hat{n}(t) = [\cos \phi(t), \sin \phi(t)]$ is the unit orientation vector pointing in the particle's direction of motion $\phi(t)$. The angular dynamics in the run phase is simply determined by a Gaussian white noise $\xi(t)$ with rotational diffusion coefficient D_r . We assume the particle to be capable of performing noisy measurements of the true target direction throughout the duration of this phase. The initial particle orientation in the n th phase is an estimate of the true target direction based on the previous run phase, which we will denote as $\hat{\theta}(\tau_{n-1})$. 2. *Inference and reorientations*: During the n th run, the particle has acquired information about the true target direction based on noisy measurements or monitoring. Once this run phase is completed, the particle uses the collected information and reorients in a direction toward the target, which becomes the initial orientation for the $(n+1)$ th phase. Concretely, let the agent spend a duration τ_n performing noisy measurements θ_j of the true navigation direction, relative to some global reference, with $j = 1, \dots, m(\tau_n)$ indexing the measurements. There are in total $m(\tau_n)$ measurements in the n th run. An estimator $\hat{\theta}(\tau_n)$ is a function of the measurements $\{\theta_j\}_{j=1}^{m(\tau_n)}$ that estimates the true target direction θ . Typically, one considers weighted linear combinations $\hat{\theta}(\tau_n) = \sum_{j=1}^{m(\tau_n)} w_j \theta_j$ for a choice of weights w_j . The precision of any unbiased estimator $[\langle \hat{\theta}(\tau_n) \rangle = \theta]$ is fundamentally limited through the Cramér-Rao inequality $\text{Var}[\hat{\theta}(\tau_n)] \geq I^{-1}(\tau_n)$, relating the estimator variance $\text{Var}[\hat{\theta}(\tau_n)]$ to the reciprocal of the *Fisher information*. For a single measurement of data θ_j coming from a distribution $p(\theta_j; \theta)$, the Fisher information is defined by $I_j = \langle [\partial_\theta \log p(\theta_j; \theta)]^2 \rangle$, where expectation is taken with respect to $p(\theta_j; \theta)$ [51,52]. For a single measurement θ_j of Gaussian data with variance σ_j^2 , one has simply $I_j = \sigma_j^{-2}$, and the total Fisher information is additive under independent measurements so that with $m(\tau_n)$ measurements in time τ_n we have $I(\tau_n) = \sum_{j=1}^{m(\tau_n)} \sigma_j^{-2}$ [51,52].

Several natural systems move akin to the above biphasic rules, for example, insects using the combination of wind direction and odors to orient [56], or dung beetles that intermittently collect information about new directions of motion from the environment or even the Milky Way [57,58]. Robotic systems can also be equipped with similar strategies [20]. The renewal structure of this model, i.e., the independence of each epoch, enables late-time dynamical properties to be obtainable from single-epoch properties, a property recently utilized in recent works to study rectified motion in the absence of information processing [59–63].




FIG. 1. Active agents use an intermittent navigation strategy that relies on information acquired during statistically independent periods to episodically reorient, generating a drift in a target direction. The n th independent epoch is characterized by an initial orientation $\hat{n}[\hat{\theta}(\tau_{n-1})]$ based on information from the previous epoch, after which the particle performs normal active Brownian motion for a time τ_n while collecting information at a rate $\hat{I} \geq 0$.

If the particle covers a mean distance $\ell_{\parallel}$ in the direction of the target in a single run phase, the effective drift speed in the target direction can be measured by $v_{\parallel} = \ell_{\parallel}/\langle\tau\rangle$, and in dimensionless form $u_{\parallel} = v_{\parallel}/v_0$. Without loss of generality, we place the target direction parallel to the x axis (see Fig. 1). By combining the active Brownian dynamics with the Cramér-Rao inequality, one arrives at a speed limit that must be obeyed for any renewal and information processing strategy:

$$u_{\parallel} \leq u_{\text{CR}} = \frac{1 - \tilde{\psi}(D_r)}{D_r \langle\tau\rangle} \int_0^{\infty} d\tau \psi(\tau) \exp\left(-\frac{1}{2I(\tau)}\right). \quad (3)$$

This constitutes the central result of our work and demonstrates general relations between particle persistence D_r^{-1} , renewal strategy [through $\psi(\tau)$ and its Laplace transform $\tilde{\psi}(s) = \int_0^{\infty} d\tau \psi(\tau) \exp(-s\tau)$], information acquisition $I(\tau)$, and a largest allowed speed $u_{\text{CR}} = u_{\text{CR}}[\psi, I]$. The speed limit depends functionally on the tuple $\{\psi, I\}$, which determines the overall strategy of the active agent. By varying these, a wide range of bounds can be derived, and generic well-established behavioral phenomena recovered, as discussed in the following. We note that the Cramér-Rao bound is saturated for so-called efficient estimators, which means that we can also interpret the bound as the speed curves for particles following the most efficient strategy. Hence, the above speed limit represents a physically attainable rather than purely theoretical bound. In the following example, estimators that saturate the inequality will be provided. For a further discussion of these matters, as well as a derivation of Eq. (3), see Appendix A.

In the following, we demonstrate the results through two natural applications, namely (I) linear information growth with constant information acquisition rate $\dot{I}$, and (II) information deterioration during internal storage. Within each information acquisition strategy, different choices of renewal durations can be explored by varying $\psi(\tau)$, revealing how particular strategies depend on stochasticity of the renewal durations.

Case I: Constant information rate—Consider the case $I(\tau) = \dot{I}\tau$, representing a situation where the particle at short (compared to τ) evenly distributed time intervals during the duration $(0, \tau)$ performs uncorrelated and identically distributed measurements of the true target direction. In this simple scenario, the optimal estimator is given by the arithmetic mean, corresponding to uniform weights $w_j = 1/m(\tau_n)$ attributed to each measurement. A derivation of this is contained by taking an appropriate limit of *Case II* discussed below. In the simplest case of deterministic periods τ , the bounding speed in Eq. (3) takes the following form:

$$u_{\text{CR}}^{\text{det}} = \frac{1 - e^{-D_r\tau}}{D_r\tau} \exp\left(-\frac{1}{2\dot{I}\tau}\right). \quad (4)$$




FIG. 2. (a) Cramér-Rao bound for the navigation speed for two different values of the information per persistence time. Solid line shows a deterministic renewal strategy [Eq. (4)], while the dashed line a stochastic one [Eq. (5)]. (b) Phase diagram where deterministic vs stochastic strategies maximize the possible navigation speed, as function of mean renewal time and information gained.

The presence of rotational diffusion $D_r > 0$ makes the drift speed decrease in time, while the presence of a positive information acquisition rate $\dot{I} > 0$ makes it grow. These competing effects lead to a nonmonotonic bound u_{CR} as a function of the run phase durations τ , as seen in Fig. 2(a). To gain analytical insights into this nonmonotonicity, we consider momentarily the regime of highly persistent active particles, where $D_r\tau \ll 1$. Expanding the first factor in the bound leads to $u_{\text{CR}} \approx (1 - D_r\tau/2) \exp[-(1/2\dot{I}\tau)]$, which has a maximum at the optimal renewal duration $\tau_{\text{opt}} = (1/4\dot{I})[\sqrt{(16\dot{I}/D_r) + 1} - 1]$. The optimal renewal duration decays monotonically with $\dot{I}$, and at large $\dot{I}$ scales as $\dot{I}^{-1/2}$. We also see that increased noise, as measured by the rotational diffusion coefficient, will decrease the optimal sensing time, reflecting the negative effect of longer periods where noise is allowed to act. Particles that are more efficient in acquiring information should act upon this information more rapidly, as this avoids the negative effects of rotational diffusion.

A relevant question in renewal-based search strategies in general is whether stochasticity in the renewal periods can enhance search [64,65]. To understand the role played by fluctuations in the sensing period, we consider Poissonian renewal times, in which case we have an exponential distribution $\psi(\tau) = re^{-r\tau}$ characterized by the rate r . Using integral representations of modified Bessel function of the second kind $K_1(z)$ [66], we find from Eq. (3) the speed limit,

$$u_{\text{CR}}^{\text{Pois.}} = \frac{r}{r + D_r} \sqrt{\frac{2r}{\dot{I}}} K_1\left(\sqrt{\frac{2r}{\dot{I}}}\right). \quad (5)$$

As in the deterministic case, the drift speed crosses over from increasing at small rates, to decaying at large rates, implying the existence of an optimal value of the reorientation rate r . The speed limits based on Poissonian

renewal times are shown in dashed lines in Fig. 2(a). We see that, in both the deterministic and stochastic case, a trade-off takes place between acting rapidly in order to avoid the directional decorrelating effect of rotational noise and having longer sensing durations that more accurately identify the target {i.e., smaller variance $\text{Var}[\hat{\theta}(\tau_j)]$ after sensing periods}. Within behavioral science, the tendency for target accuracy to grow with sensing duration reflects a speed-accuracy trade-off and is a feature found in several living systems across scales [41–47].

While the stochasticity of the renewal duration does not affect the presence of a speed-accuracy trade-off, it is interesting to observe that in certain regimes variability can either be helpful or detrimental. By considering the two cases of deterministic and Poissonian durations at the same mean $\tau = 1/r$, a fair comparison can be made where only the fluctuations in τ in the Poissonian case differ. Figure 2(b) shows regimes where the deterministic strategy is more beneficial. Generally, for small (mean) sensing periods, the particle will in the deterministic case not have collected much information, and fluctuations could help in this regard. At large mean sensing periods, fluctuations in the exponential renewal time tend to shorten the sensing duration for a substantial fraction of realizations because most of the probability mass of an exponential distribution lies below its mean. This bias toward shorter sensing times in an ensemble leads to inefficient strategies.

Case II: Deteriorating information—The framework derived can be used to study sensing strategies beyond those where information simply grows linearly. Internal noise can significantly alter the performance of intelligent agents, as demonstrated for example in bacteria [67]. If the active agent collects data as before, the data must be saved internally until the renewal period ends and the time comes for the particle to act on the information. Assuming that the internal memory storage is imperfect, the data may be corrupted or deteriorate over time. This raises an intriguing question regarding the dynamical trade-off observed above, which as we have seen, balances effects originating in noise and reorientation accuracy. How does the introduction of internal noise that acts on data enter into this trade-off?

We consider a sequence of measurements θ_j taken at evenly spaced times t_j within the run phase $(0, \tau)$. The measurements are recorded with error σ^2 at the moment they are taken, representing limits to the particle's sensory equipment. The stored data gradually deteriorates before it is used for inference at the end of the run phase. This degradation is modeled as an internal diffusion process with diffusivity D_d , independent and uncorrelated with other processes in the model, which we interpret as a deterioration rate with dimensions of inverse time. As a result, each data point accumulates an additional uncertainty that grows linearly with the time it spends in storage: beyond the intrinsic sensing noise σ^2 , a measurement taken

at time t_j acquires an extra variance proportional to the remaining time $\tau - t_j$. For convenience, we denote the incremental deterioration occurring between successive measurements by the dimensionless number $\mathcal{D} = D_d \Delta t$. The inference problem, therefore, becomes one involving heteroskedasticity, meaning the total variance is no longer uniform across the data sequence at the time of inference. Focusing on the regime of very frequent sampling, where Δt is small, one obtains the following expression for the Fisher information:

$$I(\tau) = \frac{1}{2\mathcal{D}} \log(1 + 2\mathcal{D}\dot{I}\tau), \quad (6)$$

where we introduced $\dot{I} = 1/(\sigma^2 \Delta t)$ as the information rate so that at early times $I(\tau) \approx \dot{I}\tau$. This follows from the additivity of the Fisher information under independent measurements, and a careful derivation can be found in Appendix B. We see that deterioration leads to ultraslow logarithmic growth of the Fisher information as a function of τ . The deteriorationless limit $\mathcal{D} \rightarrow 0$ recovers the previous case, which is readily verified by a simple Taylor expansion of the above logarithm. The slow logarithmic growth indicates that, due to deterioration of the data, the active agent is subject to diminishing returns whereby longer τ gives new and accurate data points but also lets old data become irrelevant. In order for the particle to reach the bound set by the Cramér-Rao inequality in this case, it must correctly weight the measurements when estimating the target direction. Given a discrete set of n_τ measurements $\{\theta_j\}$ in time τ , the optimal estimator is given by $\hat{\theta}_{\text{opt}} = I^{-1}(\tau) \sum_{j=1}^{n_\tau} [\sigma^2 + 2D_d(\tau - t_j)]^{-1} \theta_j$, as shown in Appendix C. Hence, to reach the bound, the particle must not only know all the measurement times but also the error of its sensory apparatus. This may make the bound harder to reach, although approaching the bound may be possible through evolutionary adaptation or other forms of learning.

From our general results, Eq. (3), the above Fisher information leads to a speed limit which for deterministic sensing and run phases reads

$$u_{\text{CR}}^{\text{det}}(\mathcal{D}) = \frac{1 - e^{-D_d \tau}}{D_r \tau} \exp\left(-\frac{\mathcal{D}}{\log(1 + 2\mathcal{D}\dot{I}\tau)}\right). \quad (7)$$

Note that, since $\log(1 + x) \leq x$ for all positive x , we have $I(\tau) \leq \dot{I}\tau$. It, hence, follows that also $\exp\{-1/[2I(\tau)]\} \leq \exp[-1/(2\dot{I}\tau)]$, meaning that deterioration always suppresses the speed for all values of τ . This is seen in Fig. 3(a), where examples of the speed limits are shown for different deterioration rates. Figure 3(b) shows the regions where deterministic or stochastic strategies are beneficial, where we see that internal deterioration has a small effect. The optimal renewal time can be extracted from the analytical formulas for the speed limit, resulting in




FIG. 3. (a) Cramér-Rao bound for the navigation speed without ($\mathcal{D}=0$) and with ($\mathcal{D}=5$) deterioration of the data before inference. Solid lines show deterministic renewal strategies [Eq. (7)], while the dashed line corresponds to Poissonian renewal times [obtained by numerical integration of Eq. (3) using Eq. (6)]. The dimensionless information rate is set to $\dot{I}/D_r = 5$. (b) Crossovers between optimal strategies. (c) Optimal renewal time τ_{opt} , obtained numerically from the exact analytical plots, for deterministic strategies. Dot-dashed line shows a line decaying as a square root in the information rate as a guide for the eye. (d) Optimal renewal time as a function of the dimensionless deterioration rate, with $\dot{I}/D_r = 5$.

Fig. 3(c). Here we see that, while in the deteriorationless case we have the scaling $\tau_{opt} \sim \dot{I}^{-1/2}$, as predicted from analytics in *Case I*, the presence of deterioration generally affects this scaling, making the decay slower than a square root. This reflects the need for longer durations of information collection when deterioration is present. Finally, we consider the way deterioration affects the optimum. To make analytical progress, we assume that, since large values of $\mathcal{D}$ would make any strategy slow, the realistic regime is that of relatively small values of $\mathcal{D}$. In this regime, we can expand $u_{CR}^{det}(\mathcal{D})$ to leading order in $\mathcal{D}$ and arrive at $u_{CR}^{det}(\mathcal{D}) \approx u_{CR}^{det}(0)(1 - \mathcal{D}/2)$. Since the leading order effect does not depend on τ , the same holds for stochastic strategies $u_{CR}^{Pois.}(\mathcal{D}) \approx u_{CR}^{Pois.}(0)(1 - \mathcal{D}/2)$. This shows that the leading order effect of deterioration is simply to rescale the speed limit while leaving the optimal renewal duration τ_{opt} unchanged. Figure 3(d) shows the explicit dependence of the optimum on deterioration rate, extracted numerically from the analytical formulas, where we see that, even in regimes where the above perturbation analysis should not be valid, the optimum surprisingly does not change much. This demonstrates that deterioration of internally saved data through internal noise, while lowering the speed limit, does not significantly alter the speed-accuracy trade-off, which still seems to focus on a balance of *external* noise and

reorientation accuracy. Finally, as shown in Fig. 3(b), we see that the transition into the regime where deterministic strategies dominate also remains relatively untouched once deterioration is included, but it is slightly shifted to larger values of the mean renewal time. In Fig. 3, we used Eq. (3) together with the Fisher information given by Eq. (6) and numerically integrated over exponential renewal times to obtain the speed limit in the stochastic scenarios. This shows that, while internal noise that deteriorates data can strongly suppress navigation speed, the optimal sensing duration that determines the speed-accuracy trade-off as well as the location of crossovers to regimes where fluctuating sensing periods may be beneficial remains sensitive mostly to the external dynamical noise.

Conclusions—We have analyzed a minimal model of active navigation that combines the renewal structure of intermittent active motion with information processing, revealing a speed-accuracy trade-off and optimal information acquisition strategies. We have derived an information-limited bound on the drift speed toward a stationary target direction, based on statistical inference performed by the particle. The bound, reachable through optimal strategies, is valid for a wide range of stochastic renewal durations and information processing strategies and provides analytical insights into the combined effect of persistent active motion and information processing.

Many future directions of research could be viable. Extensions that incorporate thermodynamics and enable trade-offs between accuracy, information as well as energy expenditure would be of interest [68–71]. Despite challenges in formulating a consistent thermodynamics of active systems from dynamics alone [72,73], one may still expect that both the acquisition and processing of information by intelligent active particles entail additional energetic costs [74]. In particular, information must be encoded, updated, and eventually erased, implying Landauer-type costs [75] that favor internal representations that are compressed yet sufficiently informative [76].

Furthermore, the intermittent strategy considered here could be extended to non-Markovian settings. This could arise, for example, through active agents that retain long-term memory rather than resetting it at each renewal cycle, or through correlations between renewal epochs. Such extensions naturally raise the question of whether these memory effects could further enhance navigational efficiency. Identifying optimal renewal epochs, in the form of optimal distributions $\psi(\tau)$, also represents intriguing grounds for future work. Egocentric strategies, where the particle measures target direction relative to its own body axis rather than in a global sense, would also give rise to interesting effects of correlations since multiple measurements in this case would no longer be independent.

Finally, while we here have presented results based on the mean drift speed, a better understanding of higher-order fluctuations would be insightful. For example, it may be of

interest for the particle to have reliable low-fluctuating strategies even if this comes at the cost of reduced mean navigation speed, which raises interesting further optimization questions.

Acknowledgments—K. S. O is grateful to Benjamin Friedrich and Antonio Celani for insightful discussions. K. S. O acknowledges support from the Alexander von Humboldt Foundation through their fellowship program. K. S. O and H. L acknowledge support by the Deutsche Forschungsgemeinschaft within the Project No. 572589640. M. T acknowledges financial support by JSPS KAKENHI Grant No. JP24H01485 and by JSPS International Joint Research Program (JRPs with DFG) Number JPJSJRP20261608.

Data availability—There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

-
- [1] E. Schrödinger, *What is Life? The Physical Aspect of the Living cell* (Rare Treasure Editions, Victoria, BC Canada, 2025).
 - [2] M. te Vrugt, *Artificial Intelligence and Intelligent Matter: Nanoscience, Soft Matter, Philosophy*, Machine Intelligence for Materials Science (Springer, New York, 2026).
 - [3] V. A. Baulin, A. Giacometti, D. A. Fedosov, S. Ebbens, N. R. Varela-Rosales, N. Feliu, M. Chowdhury, M. Hu, R. Fuchsli, M. Dijkstra *et al.*, Intelligent soft matter: Towards embodied intelligence, *Soft Matter* **21**, 4129 (2025).
 - [4] C. Kaspar, B. J. Ravoo, W. G. van der Wiel, S. V. Wegner, and W. H. Pernice, The rise of intelligent matter, *Nature (London)* **594**, 345 (2021).
 - [5] S. Ramaswamy, The mechanics and statistics of active matter, *Annu. Rev. Condens. Matter Phys.* **1**, 323 (2010).
 - [6] G. Tkacik and W. Bialek, Information processing in living systems, *Annu. Rev. Condens. Matter Phys.* **7**, 89 (2016).
 - [7] S. Toyabe, T. Watanabe-Nakayama, T. Okamoto, S. Kudo, and E. Muneyuki, Thermodynamic efficiency and mechanochemical coupling of F1-ATPase, *Proc. Natl. Acad. Sci. U.S.A.* **108**, 17951 (2011).
 - [8] T. Ariga, K. Tateishi, M. Tomishige, and D. Mizuno, Noise-induced acceleration of single molecule kinesin-1, *Phys. Rev. Lett.* **127**, 178101 (2021).
 - [9] G. M. Viswanathan, M. G. Da Luz, E. P. Raposo, and H. E. Stanley, *The Physics of Foraging: An Introduction to Random Searches and Biological Encounters* (Cambridge University Press, Cambridge, England, 2011).
 - [10] G. Volpe and G. Volpe, The topography of the environment alters the optimal search strategy for active particles, *Proc. Natl. Acad. Sci. U.S.A.* **114**, 11350 (2017).
 - [11] H. C. Berg and E. M. Purcell, Physics of chemoreception, *Biophys. J.* **20**, 193 (1977).
 - [12] M. J. Schnitzer, Theory of continuum random walks and application to chemotaxis, *Phys. Rev. E* **48**, 2553 (1993).
 - [13] P.-G. De Gennes, Chemotaxis: The role of internal delays, *Eur. Biophys. J.* **33**, 691 (2004).
 - [14] A. C. Tsang, E. Demir, Y. Ding, and O. S. Pak, Roads to smart artificial microswimmers, *Adv. Intell. Syst.* **2**, 1900137 (2020).
 - [15] S. Goh, R. G. Winkler, and G. Gompper, Noisy pursuit and pattern formation of self-steering active particles, *New J. Phys.* **24**, 093039 (2022).
 - [16] R. S. Negi, R. G. Winkler, and G. Gompper, Binary mixtures of intelligent active Brownian particles with visual perception, *New J. Phys.* **27**, 103301 (2025).
 - [17] J. Rode, M. Novak, and B. M. Friedrich, Information theory of chemotactic agents using both spatial and temporal gradient sensing, *PRX Life* **2**, 023012 (2024).
 - [18] L. Cocconi, B. Mahault, and L. Piro, Dissipation-accuracy tradeoffs in autonomous control of smart active matter, *New J. Phys.* **27**, 013002 (2025).
 - [19] Z. Hou, Z. Zhang, J. Li, K. Yasuda, and S. Komura, Ornstein-uhlenbeck information swimmers with external and internal feedback controls, *Europhys. Lett.* **150**, 27001 (2025).
 - [20] A. Sinha, S. Jangid, T. Sadhu, and S. Ghosh, Optimal navigation in a noisy environment, [arXiv:2512.20336](https://arxiv.org/abs/2512.20336).
 - [21] F. Mori and L. Mahadevan, Optimal switching strategies for navigation in stochastic settings, *J. R. Soc. Interface* **22**, 20240677 (2025).
 - [22] G. D. V. Del Vecchio, M. Kulkarni, S. N. Majumdar, and S. Sabhapandit, Proxitaxis: An adaptive search strategy based on proximity and stochastic resetting, *Phys. Rev. E* **113**, L042101 (2026).
 - [23] X. Song, X. Shao, Y. Zhu, C. Yang, L. He, S. Komura, and Z. Hou, Ornstein-uhlenbeck information particle: A new candidate of active agent, [arXiv:2602.06340](https://arxiv.org/abs/2602.06340).
 - [24] J. M. Betancourt, M. P. Leighton, T. Emonet, B. B. Machta, and M. C. Abbott, Discrete turn strategies emerge in information-limited navigation, [arXiv:2602.23324](https://arxiv.org/abs/2602.23324).
 - [25] S. Colabrese, K. Gustavsson, A. Celani, and L. Biferale, Flow navigation by smart microswimmers via reinforcement learning, *Phys. Rev. Lett.* **118**, 158004 (2017).
 - [26] T. Jacob, S. Mohapatra, R. A. S. Mathew, and P. Sinha Mahapatra, Mixing of a binary passive particle system using smart active particles, *Sci. Rep.* **16**, 3174 (2025).
 - [27] M. Gassner, S. Goh, G. Gompper, and R. G. Winkler, Noisy pursuit by a self-steering active particle in confinement, *Europhys. Lett.* **142**, 21002 (2023).
 - [28] M. Putzke and H. Stark, Optimal navigation of a smart active particle: Directional and distance sensing, *Eur. Phys. J. E* **46**, 48 (2023).
 - [29] F. Cichos, K. Gustavsson, B. Mehlig, and G. Volpe, Machine learning for active matter, *Nat. Mach. Intell.* **2**, 94 (2020).
 - [30] K. Gustavsson, L. Biferale, A. Celani, and S. Colabrese, Finding efficient swimming strategies in a three-dimensional chaotic flow by reinforcement learning, *Eur. Phys. J. E* **40**, 1 (2017).
 - [31] E. Schneider and H. Stark, Optimal steering of a smart active particle, *Europhys. Lett.* **127**, 64003 (2019).
 - [32] M. Durve, F. Peruani, and A. Celani, Learning to flock through reinforcement, *Phys. Rev. E* **102**, 012601 (2020).

- [33] M. Nasiri, H. Löwen, and B. Liebchen, Optimal active particle navigation meets machine learning (a), *Europhys. Lett.* **142**, 17001 (2023).
- [34] M. Nasiri, E. Loran, and B. Liebchen, Smart active particles learn and transcend bacterial foraging strategies, *Proc. Natl. Acad. Sci. U.S.A.* **121**, e2317618121 (2024).
- [35] R. A. Heinonen, L. Biferale, A. Celani, and M. Vergassola, Optimal trajectories for bayesian olfactory search in turbulent flows: The low information limit and beyond, *Phys. Rev. Fluids* **10**, 044601 (2025).
- [36] R. Singh, P. Jena, A. Kumar, and S. Mishra, Homing through reinforcement learning, *arXiv:2602.08566*.
- [37] M. Müller and R. Wehner, Path integration in desert ants, *Cataglyphis fortis*, *Proc. Natl. Acad. Sci. U.S.A.* **85**, 5287 (1988).
- [38] G. Hartmann and R. Wehner, The ant's path integration system: A neural architecture, *Biol. Cybern.* **73**, 483 (1995).
- [39] F. Lenz, T. C. Ings, L. Chittka, A. V. Chechkin, and R. Klages, Spatiotemporal dynamics of bumblebees foraging under predation risk, *Phys. Rev. Lett.* **108**, 098103 (2012).
- [40] R. Wiltchko and W. Wiltchko, Animal navigation: How animals use environmental factors to find their way, *Eur. Phys. J. Special Topics* **232**, 237 (2023).
- [41] H. E. Garrett, *A Study of the Relation of Accuracy to Speed* (Columbia university, New York, 1922), Vol. 8.
- [42] R. Ratcliff and G. McKoon, The diffusion decision model: Theory and data for two-choice decision tasks, *Neural Comput.* **20**, 873 (2008).
- [43] R. Ratcliff, P. L. Smith, S. D. Brown, and G. McKoon, Diffusion decision model: Current issues and history, *Trends Cognit. Sci.* **20**, 260 (2016).
- [44] E. D. Siggia and M. Vergassola, Decisions on the fly in cellular sensory systems, *Proc. Natl. Acad. Sci. U.S.A.* **110**, E3704 (2013).
- [45] L. Chittka, A. G. Dyer, F. Bock, and A. Dornhaus, Bees trade off foraging speed for accuracy, *Nature (London)* **424**, 388 (2003).
- [46] A. Durmaz, Y. Sarmiento, G. Fortunato, D. Das, M. E. Diamond, D. Bueti, and É. Roldán, Human perceptual decision making of nonequilibrium fluctuations, *arXiv:2311.12692*.
- [47] B. W. Brunton, M. M. Botvinick, and C. D. Brody, Rats and humans can optimally accumulate evidence for decision-making, *Science* **340**, 95 (2013).
- [48] M. Vergassola, E. Villermaux, and B. I. Shraiman, 'Info-taxis' as a strategy for searching without gradients, *Nature (London)* **445**, 406 (2007).
- [49] H. H. Mattingly, K. Kamino, B. B. Machta, and T. Emonet, *Escherichia coli* chemotaxis is information limited, *Nat. Phys.* **17**, 1426 (2021).
- [50] A. Celani, E. Villermaux, and M. Vergassola, Odor landscapes in turbulent environments, *Phys. Rev. X* **4**, 041015 (2014).
- [51] T. M. Cover and J. A. Thomas, *Elements of Information Theory* (John Wiley & Sons, New York, 1991).
- [52] S. Kullback, *Information Theory and Statistics* (Courier Corporation, 1997).
- [53] Here "speed" refers to the duration spend by the agent collecting information, not physical navigation speed. Accuracy refers to the reciprocal of variance, here the angular localization of the target direction.
- [54] J. R. Howse, Richard A. Jones, A. J. Ryan, T. Gough, R. Vafabakhsh, and R. Golestanian, Self-motile colloidal particles: From directed propulsion to random walk, *Phys. Rev. Lett.* **99**, 048102 (2007).
- [55] M. Te Vrugt and R. Wittkowski, Metareview: A survey of active matter reviews, *Eur. Phys. J. E* **48**, 12 (2025).
- [56] R. T. Cardé and M. A. Willis, Navigational strategies used by insects to find distant, wind-borne sources of odor, *J. Chem. Ecol.* **34**, 854 (2008).
- [57] M. Byrne, M. Dacke, P. Nordström, C. Scholtz, and E. Warrant, Visual cues used by ball-rolling dung beetles for orientation, *J. Comp. Physiol. A* **189**, 411 (2003).
- [58] M. Dacke, E. Baird, M. Byrne, C. H. Scholtz, and E. J. Warrant, Dung beetles use the milky way for orientation, *Curr. Biol.* **23**, 298 (2013).
- [59] V. Kumar, O. Sadekar, and U. Basu, Active brownian motion in two dimensions under stochastic resetting, *Phys. Rev. E* **102**, 052129 (2020).
- [60] Y. Baouche, T. Franosch, M. Meiners, and C. Kurzthaler, Active Brownian particle under stochastic orientational resetting, *New J. Phys.* **26**, 073041 (2024).
- [61] K. S. Olsen and H. Löwen, Optimal diffusion of chiral active particles with strategic reorientations, *Phys. Rev. E* **110**, 064606 (2024).
- [62] A. Shee, Steering chiral active Brownian motion via stochastic position-orientation resetting, *Phys. Rev. E* **113**, 025424 (2026).
- [63] S. Kundu, D. Mondal, A. Biswas, A. Pal, and M. Khan, Emulating microbial run-and-tumble and tactic motion by stochastically reorienting synthetic active Brownian particles, *arXiv:2509.21903*.
- [64] U. Bhat, C. De Bacco, and S. Redner, Stochastic search with poisson and deterministic resetting, *J. Stat. Mech.* (2016) P083401.
- [65] M. R. Evans and S. Ray, Stochastic resetting prevails over sharp restart for broad target distributions, *Phys. Rev. Lett.* **134**, 247102 (2025).
- [66] DLMF, NIST Digital Library of Mathematical Functions, edited by F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, B. V. Saunders, H. S. Cohl, and M. A. McClain, <https://dlmf.nist.gov/> (Release 1.2.0 of 2024-03-15).
- [67] H. H. Mattingly, K. Kamino, J. Ong, R. Kottou, T. Emonet, and B. B. Machta, *E. coli* chemosensing accuracy is not limited by stochastic molecule arrivals, *Nat. Phys.* **22**, 123 (2026).
- [68] M. D. Nguyen, P. H. Pham, K. V. Ngo, V. H. Do, S. Li, and T. V. Phan, Remark on the entropy production of adaptive run-and-tumble chemotaxis, *Physica A (Amsterdam)* **634**, 129452 (2024).
- [69] P. Sartori, L. Granger, C. F. Lee, and J. M. Horowitz, Thermodynamic costs of information processing in sensory adaptation, *PLoS Comput. Biol.* **10**, e1003974 (2014).
- [70] G. Lan, P. Sartori, S. Neumann, V. Sourjik, and Y. Tu, The energy-speed-accuracy trade-off in sensory adaptation, *Nat. Phys.* **8**, 422 (2012).
- [71] T. Welker and P. Pietzonka, Accuracy comes at a cost: Optimal localisation against a flow, *arXiv:2602.13173*.

- [72] R. Bebon, J. F. Robinson, and T. Speck, Thermodynamics of active matter: Tracking dissipation across scales, *Phys. Rev. X* **15**, 021050 (2025).
- [73] R. Bebon and T. Speck, Thermodynamics of active matter with internal degrees of freedom, *Commun. Phys.* **9**, 162 (2026).
- [74] L. Cocconi, B. Mahault, and L. Piro, Nonequilibrium energetics of sensing and actuation by a smart active particle, *arXiv:2603.15602*.
- [75] D. H. Wolpert, J. Korbel, C. W. Lynn, F. Tasnim, J. A. Grochow, G. Kardeş, J. B. Aimone, V. Balasubramanian, E. De Giuli, D. Doty *et al.*, Is stochastic thermodynamics the key to understanding the energy costs of computation?, *Proc. Natl. Acad. Sci. U.S.A.* **121**, e2321112121 (2024).
- [76] N. Tishby, F. C. Pereira, and W. Bialek, The information bottleneck method, *arXiv:physics/0004057*.

End Matter

Appendix A: Derivation of the bound [Eq. (3)]—From the active Brownian equations of motion, Eqs. (1) and (2), the displacement during the n th epoch is simply

$$\Delta x_{\parallel} = v_0 \int_0^{\tau_n} dt \cos(\hat{\theta}(\tau_{n-1}) + \sqrt{2D_r}W(t)). \quad (\text{A1})$$

where $W(t)$ is a Wiener process with zero mean and variance $\langle W^2(t) \rangle = t$. We note that the addition of translational zero-mean noise would not affect the following results. Performing the average over the Wiener process and using standard properties of trigonometric functions of Gaussian variables results in the conditional displacement

$$\langle \Delta x_{\parallel} | \hat{\theta}(\tau_{n-1}), \tau_n \rangle = v_0 \cos[\hat{\theta}(\tau_{n-1})] \frac{1 - e^{-D_r \tau_n}}{D_r}. \quad (\text{A2})$$

We emphasize that other models of active motion, such as the run and tumble model, could also be used, resulting in similar expressions but with rotational diffusion exchanged with a tumbling rate. Three random variables remain to be averaged in order to get the mean displacement: the two independent and identically distributed times $\{\tau_{n-1}, \tau_n\}$, and the estimator $\hat{\theta}$. We first perform the average over the estimator, assuming that it is Gaussian distributed with zero mean. This assumes that, for example, if measurements are Gaussian distributed, the particle acts on the collected information through some linear combination of the data. This includes strategies such as picking the median, mean, extrema, or midpoint of the data to name a few. Averaging over the estimator gives a displacement depending only on the two statistically independent renewal times,

$$\langle \Delta x_{\parallel} | \tau_{n-1}, \tau_n \rangle = v_0 \frac{1 - e^{-D_r \tau_n}}{D_r} \exp\left(-\frac{\text{Var}[\hat{\theta}(\tau_{n-1})]}{2}\right). \quad (\text{A3})$$

The Cramér-Rao bound from statistical inference tells us that the variance of this estimator is bounded by the reciprocal of the Fisher information $\text{Var}[\hat{\theta}(\tau_{n-1})] \geq$

$I^{-1}(\tau_{n-1})$. For a single measurement of data θ_j coming from a distribution $p(\theta_j; \theta)$, one has $I_j = \langle [\partial_{\theta} \log p(\theta_j; \theta)]^2 \rangle$ [51,52]. Rather than assuming an explicit model for the individual measurements made by the particle during the information acquisition phase, we will assume that the full phase leads to a total Fisher information $I(\tau)$ after many measurements. The functional dependence of $I(\tau)$ on τ defines the model's information acquisition strategy. Since the exponential is monotonically decaying, the Cramér-Rao bound can also be stated as $\exp\{-(\text{Var}[\hat{\theta}(\tau_{n-1})]/2)\} \leq \exp\{-[1/2I(\tau_{n-1})]\}$, directly offering a bound for the particle displacement in Eq. (A3). Averaging over the two remaining independent and identically distributed durations to find the displacement $\ell_{\parallel} = \langle \Delta x_{\parallel} | \tau_{n-1}, \tau_n \rangle_{n-1,n}$ results in the speed limit

$$u_{\parallel} \leq u_{\text{CR}} = \frac{1 - \tilde{\psi}(D_r)}{D_r \langle \tau \rangle} \int_0^{\infty} d\tau \psi(\tau) \exp\left(-\frac{1}{2I(\tau)}\right), \quad (\text{A4})$$

where we defined the (dimensionless) mean speed as $u_{\parallel} = v_{\parallel}/v_0 = \ell_{\parallel}/(\langle \tau \rangle v_0)$, and $\tilde{\psi}(s)$ is the Laplace transform of $\psi(\tau)$.

Appendix B: Fisher information under internal memory deterioration—Each data point has, in addition to error σ^2 coming from the measurement itself, an additional variance $2D_d(\tau - t_k)$ due to the internal deterioration taking place before inference is made. Since the Fisher information is additive for independent measurements, it is given by the sum

$$I(\tau) = \sum_{k=1}^m \frac{1}{\sigma^2 + 2D_d(\tau - t_k)} = \sum_{k=1}^m \frac{1}{\sigma^2 + 2D(n - k)}, \quad (\text{B1})$$

where we considered a discretization of size Δt in time so that $t_j = j\Delta t$, with $\tau = t_m = m\Delta t$, and used $D = D_d\Delta t$. Note that by reparametrizing the sum in terms of $j = m - k$ we have

$$I(\tau) = \frac{1}{2\mathcal{D}} \sum_{j=0}^{m-1} \frac{1}{\sigma^2/(2\mathcal{D}) + j}. \quad (\text{B2})$$

To proceed, we start by noting that repeated use of the recursive property $\Gamma(z+1) = z\Gamma(z)$ of the Gamma function leads to the identity $\Gamma(z+m) = z^{(m)}\Gamma(z)$, where $z^{(m)} = z(z+1)\cdots(z+m-1)$ is the Pochhammer ascending factorial. Taking a logarithmic derivative gives

$$\frac{d}{dz} \log \Gamma(z+m) = \sum_{j=0}^{m-1} \frac{1}{z+j} + \frac{d}{dz} \log \Gamma(z). \quad (\text{B3})$$

Rearranging and applying this identity to our above sum results in

$$I(\tau) = \frac{1}{2\mathcal{D}} \left[\frac{\Gamma'(\frac{\sigma^2}{2\mathcal{D}} + m)}{\Gamma(\frac{\sigma^2}{2\mathcal{D}} + m)} - \frac{\Gamma'(\frac{\sigma^2}{2\mathcal{D}})}{\Gamma(\frac{\sigma^2}{2\mathcal{D}})} \right]. \quad (\text{B4})$$

We note that in the $\mathcal{D} \rightarrow \infty$ limit this reduces to $I(\tau) = 1/\sigma^2$, which is the Fisher information of a single Gaussian measurement at the final time τ . In this limit, any measurement prior to the reorientation time at the end of the run phase immediately deteriorates and becomes irrelevant, and the particle only uses the most recent data point measured exactly at the time of reorientation.

While Eq. (B4) is exact and already can be used in Eq. (3) with $m = \tau/\Delta t$, we will for the sake of transparent interpretation simplify the expression with a minimal approximation. When the sampling rate is high $\Delta t \ll 1$, we can use the asymptotic expansion $\Gamma'(z)/\Gamma(z) = \log(z) + \mathcal{O}(z^{-1})$, leading to

$$I(\tau) = \frac{1}{2\mathcal{D}} \log(1 + 2\mathcal{D}\dot{I}\tau), \quad (\text{B5})$$

where $\dot{I} = (\sigma^2\Delta t)^{-1}$, which is used and discussed in the main text.

Appendix C: Maximal likelihood strategies—The speed limit set by the Cramér-Rao bound is tight, in the sense that known strategies can reach it. Indeed, for the case of independent and identically distributed measurements, the optimal estimate the agent can use is simply the arithmetic mean. For the case with deterioration, the optimal strategy is more complex and depends on the time of the measurements. The optimal estimator is often the maximum likelihood estimator (MLE), obtained by solving $\partial_{\hat{\theta}_{\text{MLE}}} \log P(\{\theta_j\}; \hat{\theta}_{\text{MLE}}) = 0$, where P denotes the joint distribution of the data. Since each measurement is independent, $P(\{\theta_j\}; \theta) = \prod_j p(\theta_j; \theta)$ with $p(\theta_j; \theta) \sim \mathcal{N}[\theta, \sigma^2 + 2\mathcal{D}_d(\tau - t_j)]$. The

MLE is, therefore, obtained by solving

$$\sum_{j=1}^{m(\tau)} \frac{\theta_j - \hat{\theta}_{\text{MLE}}}{\sigma^2 + 2\mathcal{D}_d(\tau - t_j)} = 0. \quad (\text{C1})$$

The solution is given by

$$\hat{\theta}_{\text{MLE}} = \frac{1}{I(\tau)} \sum_{j=1}^{m(\tau)} \frac{\theta_j}{\sigma^2 + 2\mathcal{D}_d(\tau - t_j)}, \quad (\text{C2})$$

where $m(\tau)$ is the number of measurements taken in the period τ and we recognize the Fisher information as the sum of inverse variances. This one readily verifies that $\text{Var}(\hat{\theta}_{\text{MLE}}) = 1/I(\tau)$. Hence, the optimal strategy where $u_{\parallel} = u_{\text{CR}}$ must consider a weighted sum of the data, where old and more deteriorated measurements are given less weight.

Appendix D: Remarks on finite-time reorientations—In the model discussed in the main text, the active agent reorients instantaneously once the time comes to act on the collected information. This corresponds to situations where the reorientations are sufficiently fast compared to the other timescales in the system. The effect of finite-time reorientations has been studied, for example, in Ref. [20], where it is shown that it may give rise to optimal reorientation frequencies. To see how finite-time reorientations affect the present model, we consider for simplicity the case of deterministic sensing periods. While the distance covered by the agent in one sensing epoch remains the same, the time over which it takes place is extended by a reorientation time τ_R . Hence, the navigation speed is $u_{\parallel}^{\text{F.T}} = \ell_{\parallel}/(\tau + \tau_R)$, which equals the speed considered in the main text rescaled by the fraction of time spent reorienting. When rescaling functions by this fraction $u_{\parallel}^{\text{F.T}} = u_{\parallel}(1 + \tau_R/\tau)^{-1}$, the sign of the slope of $u_{\parallel}^{\text{F.T}}$ evaluated at the optimal sensing time of $u_{\parallel}$ is always positive, i.e.,

$$\text{sign} \left[\frac{du_{\parallel}^{\text{F.T}}(\tau)}{d\tau} \Big|_{\tau_*} \right] = \text{sign}[u_{\parallel}(\tau_*)], \quad (\text{D1})$$

which follows simply from the product and chain rule of differentiation. Hence, since $u_{\parallel}(\tau_*) \geq 0$, the slope of the speed including finite-time reorientations is positive, meaning it has not yet reached its maximum. Hence, the optimal sensing duration shifts in the positive direction once finite-time reorientations are included. This reflects the fact that shorter sensing durations are no longer beneficial since they now come at the additional penalty of more frequently not moving.