

Optimal Closed-Loop Control of Active Particles and a Minimal Information Engine

Rosalba Garcia-Millan^{1,2,3,*} Janik Schüttler^{2,*} Michael E. Cates² and Sarah A. M. Loos^{2,†}

¹*Department of Mathematics, King's College London, London WC2R 2LS, United Kingdom*

²*DAMTP, Centre for Mathematical Sciences, University of Cambridge, Cambridge CB3 0WA, United Kingdom*

³*St John's College, Cambridge CB2 1TP, United Kingdom*



(Received 24 January 2025; revised 27 April 2025; accepted 8 July 2025; published 18 August 2025)

We study the elementary problem of moving an active particle by a trap with minimum work input. We show analytically that (open-loop) optimal protocols are not affected by activity, but work fluctuations are always increased. For closed-loop protocols, which rely on initial measurements of the self-propulsion, the average work has a minimum for a finite persistence time. Using these insights, we derive an optimal periodic active information engine, which is found to have higher precision and information efficiency when operated with a run-and-tumble particle than for an active Ornstein-Uhlenbeck particle and, we argue, than for any other type of active particle.

DOI: [10.1103/physrevlett.135.088301](https://doi.org/10.1103/physrevlett.135.088301)

Exploring strategies to control the individual or collective motion of self-propelled particles [1–5] enhances our understanding of active dynamics and represents a crucial step toward technological applications, ranging from active metamaterials to nanosized robots. Recent research has investigated various means of control, such as adjusting the potential landscape [6–12], confinement [13–15], or the level of intrinsic activity [2]. In the realm of small-scale organisms and machines, high frictional losses and the prevalence of (non)thermal fluctuations become particularly important. In this context, “thermodynamically optimal” control schemes minimize the heat dissipation, work input, or entropy production associated with transitioning a system from one state to another. Even for passive systems, such schemes can lead to nonintuitive solutions with driving discontinuities [16–20]. To explore the principles of the thermodynamically optimal control of active systems, which are characterized by noise, memory, (non-) Gaussian statistics, and intrinsic driving, it is thus vital to study paradigmatic simplified models.

In this spirit, we investigate the elementary problem of dragging a harmonic trap containing a single active particle to a target position within a given finite time, through some fluid forming a heat bath, in one spatial dimension. We optimize the dragging protocol to minimize the

experimenter’s work input, a directly measurable quantity [20,21] that bounds below the controller’s total energy expenditure. The problem is complex enough to show key features of optimal control of active matter, yet can be solved exactly. Different from passive particles (PPs), for which this problem was solved in the seminal paper [16], self-propelled particles can actively swim toward the target, suggesting that less external work input is needed. However, the internal self-propulsion can just as well increase the resistance against the translation. We explore under which conditions the intrinsic activity facilitates or hinders the transport. To gain broader insights, we consider two models: run-and-tumble particles (RTPs) and active Ornstein-Uhlenbeck particles (AOPs).

From a control theory perspective, the dragging problem represents an *open-loop* control scheme, because the protocol is predefined and fixed. More sophisticated are *closed-loop* or feedback control schemes, which dynamically adapt the protocol taking into account real-time knowledge of the stochastic system state, resembling a mesoscopic “Maxwell demon” [21]. Closed-loop control can generally provide more precise manipulation, but it necessarily relies on measurements which, according to Landauer’s principle, have a thermodynamic cost [22–25]. As a simple example of closed-loop control, we study how an initial measurement of the system state affects the optimal protocol, generalizing previous studies on PPs [19,26] to the active case. Remarkably, we find that if the particle is initially moving toward the target position, the optimal protocol can involve an initial jump of the trap *away* from the target, challenging naive intuition.

As an application, we theoretically construct a minimal, thermodynamically optimized iteration that uses measurement information and returns extractable work, creating an “information engine” [22,27,28]. The harvested work

*These authors contributed equally to this work.

†Contact author: sarah.loos@ds.mpg.de

Present address: Max Planck Institute for Dynamics and Self-Organization, 37077 Göttingen, Germany.

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.

stems from the in-built activity [13,14,28,29], as in an active heat engine [30–34].

Our approach complements recent perturbative methods based on response theory [6,7], which enable the study of more complex active systems, but are limited to the weak and slow driving regime. Our exact results allow for an explicit analysis of fast and intermediate driving, where driving discontinuities have been observed in various contexts for PPs [16–20].

Model—We consider a one-dimensional model for a self-propelled particle at position x in a harmonic potential $V(x, \lambda) = k(x - \lambda)^2/2$ centered at λ with stiffness k . The particle's motion follows the overdamped Langevin equation

$$\dot{x}(t) = -\partial_x V(x, \lambda) + v(t) + \sqrt{2D}\xi(t), \quad (1)$$

where v denotes the self-propulsion, D is the thermal diffusion constant, and ξ is a thermal, unit-variance Gaussian white noise, $\langle \xi(t) \rangle = 0$, $\langle \xi(t)\xi(t') \rangle = \delta(t - t')$. We have absorbed the friction constant and k_B into other parameters. We consider two standard models of active motility: (i) AOUPs, where $\tau \dot{v}(t) = -v(t) + \sqrt{2D_v}\xi_v(t)$, with self-propulsion “diffusion” constant D_v and unit-variance Gaussian white noise ξ_v independent of ξ ; and (ii) RTPs, where v is a dichotomous Poissonian noise that randomly reassigns values $\pm\omega$ with tumbling rate $1/\tau$ [1,35–38]. A direct comparison can be established up to second-order correlators by fixing $D_v/\tau \equiv \omega^2$, such that both models yield

$$\langle v(t)v(t') \rangle_s = \omega^2 e^{-|t-t'|/\tau} \quad (2)$$

at stationarity (s), with correlation time τ . Comparing AOUPs versus RTPs allows us to study the effect of non-Gaussianity (RTPs) [39]. We further compare with PPs, where $v(t) \equiv 0$.

Both active models reduce to PPs under the limits $\omega \rightarrow 0$ or $\tau \rightarrow 0$. The limit $\tau \rightarrow \infty$ yields a ballistic motion with a constant random velocity set by the initial condition.

We study the problem of moving the trap from $\lambda_0 = 0$ at $t = 0$, where the system is in a nonequilibrium steady state, to a target position $\lambda(t_f) = \lambda_f$ in a finite time t_f . Owing to the frictional resistance of the fluid and changes in the particle's potential energy, such a process generally requires external work input by the controller [40,41],

$$W = \int_0^{t_f} dt \dot{\lambda} \frac{\partial V(x, \lambda)}{\partial \lambda} = \int_0^{t_f} dt \dot{\lambda} k(\lambda - x). \quad (3)$$

We optimize the time-dependent protocol $\lambda(t)$ to minimize the average work input. We deliberately do not account for the dissipation associated with self-propulsion, as the internal swimming mechanism is unresolved at the model level, nonuniversal, and hard to measure experimentally (to quantify it, one could explicitly model the swimming mechanism [33,42–44]).

Open-loop control—First we discuss the open-loop control problem where no measurements are taken. We minimize the noise-averaged work functional given in Eq. (3) using the Euler-Lagrange formalism; see Ref. [45] for the explicit derivation. Thereby, the initial condition is drawn from the nonequilibrium steady state, where $\langle x(0) \rangle_s = 0$ and $\langle v(0) \rangle_s = 0$. We denote averages over trajectories initialized in the steady state by $\langle \dots \rangle_s$. Interestingly, the resulting optimal protocol λ turns out to be independent of the self-propulsion, and therefore, *identical* for AOUPs, RTPs, and PPs [16]. It has two symmetric jumps at $t = 0, t_f$ of length $\lambda_f/(k t_f + 2)$ connected by a linear dragging regime, shown as a thin line in Fig. 1(b), and its energetic cost is $\langle W \rangle_s = \lambda_f^2/(t_f + 2/k) \geq 0$. Thus, the activity does not decrease the required average work. These




FIG. 1. Optimal dragging for open-loop control and for closed-loop control, where the demon measures the particle's initial self-propulsion (here, $v_0 > 0$ as an example). (a) Protocol stages (1 → 2 initial jump, 2 → 3 continuous driving regime, 3 → 4 final jump to target). (b) Protocol (blue), average particle position (gray) and self-propulsion (arrows). The thin blue line shows the open-loop protocol; the thick line shows the corresponding closed-loop protocol. (c) Cumulative average work $\langle W(t) \rangle_{v_0} = \int_0^t dt' \dot{\lambda} k(\lambda - \langle x \rangle_{v_0})$ for open-loop control (thin line) and for closed-loop control (thick line). Here, $\tau = 0.2$, $k = 1$, $t_f = 1$, and $v_0 = 1$.

findings are important in experimental settings, where the underlying particle dynamics and level of activity may be unknown.

However, the activity influences the *fluctuations* of the work done by the controller. As we derive in the companion paper [45], the variance of the work for PPs is $\text{Var}_s W_{\text{passive}} = 2D\lambda_f^2/(t_f + 2/k)$, while for both active models, there is an identical additional contribution:

$$\begin{aligned} & \text{Var}_s W_{\text{active}} - \text{Var}_s W_{\text{passive}} \\ &= \frac{2\lambda_f^2\omega^2}{\left(\frac{t_f}{\tau} + \frac{2}{k\tau}\right)^2} \left[\frac{t_f}{\tau} + \frac{\frac{2}{k\tau} + (1 - k\tau)(1 - e^{-t_f/\tau})}{1 + k\tau} \right] \geq 0. \end{aligned} \quad (4)$$

Thus, the variance of the work is always increased by activity, consistent with the picture of an increased effective temperature. The additional variance only vanishes in the trivial (passive) limits: $\lambda_f \rightarrow 0$, $\omega \rightarrow 0$, and in the quasi-static limit, $t_f \rightarrow \infty$, where $\text{Var}_s W$ vanishes for all models.

Taken together, our findings suggest that activity provides no advantage in the open-loop problem: while the optimal protocol and minimum mean work remain identical to the passive case, the stochastic work entails larger uncertainty. Furthermore, at the level of first and second moments, the (non-)Gaussianity is irrelevant for open-loop control, though a distinction arises in the higher moments. For example, the positional distribution remains Gaussian for PPs and AOUPs throughout the protocol, while the non-Gaussian distribution of RTPs is generally changed by the dragging [38].

Closed-loop control (initial measurement)—We turn to the case where a feedback controller or mesoscopic “demon” takes an initial real-time measurement before deciding upon the protocol. The optimal closed-loop protocol can be obtained by the same procedure as before, now incorporating the (measured) states $x(0) = x_0$, $v(0) = v_0$ as the initial condition. We focus on the case where only v is measured, which is obtained by marginalizing over x . Notably, a measurement of v also yields statistical information about x , since the average of $x(0)$, conditioned on $v(0) = v_0$, is $\langle x(0) \rangle_{v_0} = v_0/(k + \tau^{-1}) \neq \langle x(0) \rangle$ [38,45]; here we introduced the notation $\langle \bullet \rangle_{v_0} = \langle \bullet | v(0) = v_0 \rangle$.

As detailed in the companion paper [45], the optimal protocol after a v measurement, which again coincides for AOUPs and RTPs, reads as

$$\lambda_{v_0}(t) = \begin{cases} 0, & t = 0 \\ \langle x(t) \rangle_{v_0} + \frac{d_{v_0}}{kt_f + 2} - \frac{v_0}{2k} e^{-t/\tau} & 0 < t < t_f, \\ \lambda_f, & t = t_f \end{cases} \quad (5)$$

with

$$\langle x(t) \rangle_{v_0} = \frac{v_0}{k + \tau^{-1}} + \frac{d_{v_0}}{t_f + 2/k} t + \frac{\tau v_0}{2} (1 - e^{-t/\tau}), \quad (6)$$

$$d_{v_0} = \lambda_f - \frac{v_0}{k + \tau^{-1}} - \frac{\tau v_0}{2} (1 - e^{-t_f/\tau}). \quad (7)$$

The quantity d_{v_0} corresponds to the distance between the target trap position λ_f and the expected value of the particle’s initial position after v_0 measurement, minus the mean distance covered “for free” by the particle’s self-propulsion until its orientation decorrelates (last term).

Figures 1(a) and 1(b) show a typical optimal protocol (5) for some $v_0 > 0$, $\lambda_f > 0$. Different from the open-loop case, the jumps at $t = 0, t_f$ can be asymmetric, and, at intermediate times, the protocol is generally nonlinear. Furthermore, the initial jump

$$\lambda_{v_0}(0^+) - \lambda_{v_0}(0) = \frac{v_0}{k + \tau^{-1}} - \frac{v_0}{2k} + \frac{d_{v_0}}{kt_f + 2}, \quad (8)$$

can be in the *opposite* direction to the target position λ_f . What seems at first glance even more counterintuitive, these “reversed” jumps occur when the initial self-propulsion already pushes the particle toward λ_f ; see Fig. 1(b). The expectation value for the corresponding cumulative work, $\int_0^t dt' \dot{\lambda}(\lambda - \langle x \rangle)$, shown in Fig. 1(c), reveals the reason. While the initial reversed jump is costly, it brings the system into a configuration that enables a subsequent *extraction* of net work from the self-propulsion (negative $\langle dW \rangle$), overcompensating the initial cost.

The average total work associated with λ_{v_0} is

$$\langle W \rangle_{v_0} = -\frac{kv_0^2}{2(k + \tau^{-1})^2} - \frac{\tau v_0^2}{8} (1 - e^{-2t_f/\tau}) + \frac{d_{v_0}^2}{t_f + 2/k}, \quad (9)$$

which implies that v measurements enable work extraction through two distinct physical mechanisms. First, measuring v gives average positional information that makes it possible to extract some of the potential energy stored in the initial configuration, concretely $(k/2)\langle x(0) \rangle_{v_0}^2$ [first term in Eq. (9)] [46]. Second, energy can be extracted from the self-propulsion: setting the trap to a position where the particle actively *climbs up* the potential and then moving the trap along the direction of self-propelled motion, one can extract work from the activity until the orientation decorrelates. In contrast to work extracted from potential energy, this operation requires a finite duration, $t_f > 0$, unlike work extraction from potential energy of PPs through positional measurements [26].

Averaging Eq. (9) over the steady state distribution of $v(0)$ yields the expected mean work

$$\begin{aligned} \mathcal{W} = & -\frac{k\omega^2}{2(k + \tau^{-1})^2} - \frac{\tau\omega^2}{8} (1 - e^{-2t_f/\tau}) \\ & + \frac{\lambda_f^2 + \left(\frac{2}{k + \tau^{-1}} + \tau(1 - e^{-t_f/\tau})\right)^2 \omega^2/4}{t_f + 2/k}. \end{aligned} \quad (10)$$




FIG. 2. The work for closed-loop control averaged over v_0 measurements exhibits a pronounced minimum for a finite persistence time τ (here, $k = 1$). The y axis shows the reduction in average work compared to the passive limit ($\tau \rightarrow 0$). The optimum value τ^* (black line) is nonmonotonic in t_f , due to a “kink” around $t_f = 1$, discussed in more detail in [45].

$\mathcal{W}$ has a pronounced minimum at finite τ , Fig. 2. Thus, the expected cost to translate an active particle is lower than for a PP ($\tau \rightarrow 0$), and, remarkably, also a ballistic particle ($\tau \rightarrow \infty$). In both these limits the work $\mathcal{W}$ reduces to $(\lambda_f^2/t_f + 2/k)$. Therefore, a finite persistence time, i.e., the presence of nonequilibrium fluctuations, poses an advantage for closed-loop control.

Active information engine—A notable potential application of optimal protocols is optimized periodic machines. Based on our results so far, one can readily construct a minimal, optimal information engine that harvests energy from activity using self-propulsion measurements, as follows. From Eqs. (5)–(9), we find that the work $\langle W \rangle_{v_0}$ is

minimized, i.e., work extraction is maximized, if the trap is set to $\lambda_f^* = v_0/(k + \tau^{-1}) + \frac{1}{2}\tau v_0(1 - e^{-t_f/\tau})$. For this optimal target displacement λ_f^* , the final trap position coincides with the average particle position at the end of the protocol, so that the protocol extracts any average potential energy available at $t = t_f$. Setting $\lambda_f = \lambda_f^*$ further yields $\langle W \rangle_{v_0} \leq 0$ for all v_0 , and $d_{v_0} = 0$. The corresponding optimal protocol reads as

$$\lambda^{\text{engine}} = \begin{cases} 0, & t = 0 \\ \frac{\tau v_0}{2} \left(\frac{3+k\tau}{1+k\tau} - \frac{1+k\tau}{k\tau} e^{-t/\tau} \right) & 0 < t < t_f \\ \frac{\tau v_0}{2} \left(\frac{3+k\tau}{1+k\tau} - e^{-t_f/\tau} \right) & t = t_f \end{cases} \quad (11)$$

Iterating λ^{engine} yields an information engine with period t_f , visualized in Fig. 3. Over long time scales, the trap center undergoes diffusive motion. This can however be prevented by reset events that do not asymptotically affect efficiency [47].

The motor builds up nontrivial correlations between iterations, so the system does not reach a steady state. However, the orientational dynamics is not affected by the trap, and the trap position automatically coincides with the average particle position at the end of the protocol. Therefore, assuming that t_f is large compared to typical intrinsic relaxation times [48], we can still calculate the average work extraction per iteration from Eq. (9), which yields $(k/2)\omega^2 v_0(k + \tau^{-1})^{-2} - \frac{1}{8}\omega^2 \tau(1 - e^{-2t_f/\tau})$ for AOUPs and RTPs. This has no optimum τ , and reaches its maximum in the quasistatic regime. Simulating this engine, we find that, while the average work is identical for AOUPs and RTPs, the




FIG. 3. Active information engine: (a) iterated engine protocol determined upon an initial self-propulsion measurement v_0 in each iteration, (b) single realization of engine protocol (thick blue), corresponding noise-averaged particle position (gray), and cumulative work (thick red). The arrows indicate the measured v_0 values. Thin lines show trajectories for other stochastic realizations ($\omega = 1.0, \tau = 0.5, t_f = 10, k = 0.5$). (c) Work distribution per iteration for AOUP and RTP from simulations ($\omega = 0.2, \tau = 1, t_f = 10, k = 2$). While the mean (times symbol) is identical for AOUP and RTP, matching the theory (dashed line), AOUPs have a larger standard deviation (vertical bars). For RTPs, the unique mode of the distribution is below the mean, while it is close to zero for AOUPs.

work distributions differ significantly, and most notably are always broader for AOUPs, Fig. 3(c).

Finally, we address the effect of measurement uncertainty and the thermodynamic cost of information acquisition [22,49–51], which has not been included in previous studies of active information engines [14]. The uncertainty of the v measurement lowers the extractable work. Assuming a Gaussian-distributed measurement error ϵ , we show in [45] that the net extracted work has an additional cost term $\propto \epsilon^2$, leading to a total average work per iteration of

$$\mathcal{W}^{\text{engine}} = (\epsilon^2 - \omega^2) \left(\frac{k/2}{(k + \tau^{-1})^2} + \frac{\tau}{8} (1 - e^{-2t_f/\tau}) \right). \quad (12)$$

Thus, work extraction is still possible for all finite k, τ, t_f , provided that $\epsilon < |\omega|$.

The information efficiency η of the engine is obtained by dividing the mean extracted work per iteration by the information acquisition cost. Henceforth, we assume an isothermal process, where the demon operates at temperature $D = D_v$. Under this assumption, the fundamental limit to the information acquisition cost for an individual measurement is given by the relative entropy and, when averaged over multiple iterations, by the mutual information $\mathcal{I}$ between the true value v_0^r and the measurement outcome v_0^m , scaled by the temperature [26,49,51–53]. Thus,

$$\eta = \frac{-\mathcal{W}^{\text{engine}}}{D_v \mathcal{I}}. \quad (13)$$

For AOUPs, the mutual information is [45]

$$\mathcal{I}_{\text{AOUP}}[v_0^m; v_0^r] = \frac{1}{2} \ln \left(1 + \frac{\omega^2}{\epsilon^2} \right), \quad (14)$$

which vanishes for $\epsilon \rightarrow \infty$, where measurements become useless, and diverges for $\epsilon \rightarrow 0$. This divergence is not surprising, in view of the infinite information content in a continuous variable, $v_0 \in \mathbb{R}$ [54,55]. For RTPs, we find in the regime of small errors $\epsilon \ll \omega$,

$$\mathcal{I}_{\text{RTP}}[v_0^m; v_0^r] = \ln(2), \quad (15)$$

i.e., the measurement yields exactly *one bit* of information. Comparing (14) and (15) indicates that, at least in the regime of interest of modest ϵ , the cost of information acquisition is generally lower for RTPs than AOUPs. In fact, RTPs have the smallest information cost (one bit) of *any* possible active particle model. Recalling that the extractable work is independent of the model of self-propulsion, we conclude that the information efficiency η is highest for RTPs.

Together with Eqs. (12) and (15), this gives a universal upper bound to the information efficiency,

$$\eta \leq \eta_{\text{max}} = 1/[4 \ln(2)]. \quad (16)$$

The engine is most efficient when run with an RTP at low frequency $1/t_f$, with low error ϵ , and $\tau = k^{-1}$ [56]. Even for finite t_f , the efficiency η admits a maximum at a finite persistence time τ .

Conclusions—We have generalized a canonical optimization problem to the case of active particles. The open-loop protocol is not affected by activity, but fluctuations of the work are generally increased. In contrast, closed-loop protocols can extract energy from self-propulsion measurements by jumps away from the target position. Surprisingly, a finite persistence time is beneficial compared to “infinitely persistent” (ballistic) particles. We have further derived an optimal engine to harvest work from the self-propulsion measurements by simple shifts of the trap. Such an engine is most efficient and most precise for RTPs benefiting from non-Gaussianity. We found a universal upper information efficiency bound for this engine design of $1/[4 \ln(2)] \approx 0.36$. Higher efficiencies are presumably achievable with more complex machines, e.g., by a dynamic adjustment of trap stiffness.

Our results for AOUPs readily generalize to higher dimensions, and it would be compelling to explore other active models in higher dimensions too. To address more complicated processes, recent machine learning algorithms offer a promising approach [18–20], potentially enabling the study of collective active systems. We hope that our work will stimulate experimental investigations into optimal control of optically trapped bacteria [57] or synthetic active particles [58,59].

Acknowledgments—We thank Robert Jack and Édgar Roldán for insightful discussions. R. G.-M. acknowledges support from a St John’s College Research Fellowship, Cambridge. J. S. acknowledges funding through the UK Engineering and Physical Sciences Research Council (Grant No. 2602536). S. A. M. L. acknowledges funding by UK Research and Innovation (UKRI) under the UK government’s Horizon Europe funding Guarantee (Grant No. EP/X031926/1), through the Walter Benjamin Stipendium (Project No. 498288081) from the Deutsche Forschungsgemeinschaft (DFG), and from Corpus Christi College, Cambridge.

Data availability—The data that support the findings of this Letter are openly available [60].

-
- [1] Elias Schneider and Holger Stark, Optimal steering of a smart active particle, *Europhys. Lett.* **127**, 64003 (2019).
 - [2] Martin J Falk, Vahid Alizadehyazdi, Heinrich Jaeger, and Arvind Murugan, Learning to control active matter, *Phys. Rev. Res.* **3**, 033291 (2021).
 - [3] Benno Liebchen and Hartmut Löwen, Optimal navigation strategies for active particles, *Europhys. Lett.* **127**, 34003 (2019).

- [4] Marco Baldovin, David Guéry-Odelin, and Emmanuel Trizac, Control of active Brownian particles: An exact solution, *Phys. Rev. Lett.* **131**, 118302 (2023).
- [5] Suraj Shankar, Vidya Raju, and L. Mahadevan, Optimal transport and control of active drops, *Proc. Natl. Acad. Sci. U.S.A.* **119**, e2121985119 (2022).
- [6] Luke K Davis, Karel Proesmans, and Étienne Fodor, Active matter under control: Insights from response theory, *Phys. Rev. X* **14**, 011012 (2024).
- [7] Deepak Gupta, Sabine H. L. Klapp, and David A. Sivak, Efficient control protocols for an active Ornstein-Uhlenbeck particle, *Phys. Rev. E* **108**, 024117 (2023).
- [8] Erik Aurell, Carlos Mejía-Monasterio, and Paolo Muratore-Ginanneschi, Optimal protocols and optimal transport in stochastic thermodynamics, *Phys. Rev. Lett.* **106**, 250601 (2011).
- [9] Timothy Ekeh, Michael E. Cates, and Étienne Fodor, Thermodynamic cycles with active matter, *Phys. Rev. E* **102**, 010101(R) (2020).
- [10] Rita Majumdar, Arnab Saha, and Rahul Marathe, Exactly solvable model of a passive Brownian heat engine and its comparison with active engines, *J. Stat. Mech.* (2022) P073206.
- [11] Steven Blaber and David A. Sivak, Optimal control in stochastic thermodynamics, *J. Phys. Comm.* **7**, 033001 (2023).
- [12] Hyun Keun Lee, Youngchae Kwon, and Chulan Kwon, Nonequilibrium thermodynamics for a harmonic potential moving in time, *J. Korean Phys. Soc.* **83**, 331 (2023).
- [13] Paolo Malgaretti and Holger Stark, Szilard engines and information-based work extraction for active systems, *Phys. Rev. Lett.* **129**, 228005 (2022).
- [14] Luca Cocconi, Jacob Knight, and Connor Roberts, Optimal power extraction from active particles with hidden states, *Phys. Rev. Lett.* **131**, 188301 (2023).
- [15] Corneel Casert and Stephen Whitelam, Learning protocols for the fast and efficient control of active matter, *Nat. Commun.* **15**, 9128 (2024).
- [16] Tim Schmiedl and Udo Seifert, Optimal finite-time processes in stochastic thermodynamics, *Phys. Rev. Lett.* **98**, 108301 (2007).
- [17] Steven Blaber, Miranda D. Louwerse, and David A. Sivak, Steps minimize dissipation in rapidly driven stochastic systems, *Phys. Rev. E* **104**, L022101 (2021).
- [18] Megan C. Engel, Jamie A. Smith, and Michael P. Brenner, Optimal control of nonequilibrium systems through automatic differentiation, *Phys. Rev. X* **13**, 041032 (2023).
- [19] Stephen Whitelam, Demon in the machine: Learning to extract work and absorb entropy from fluctuating nano-systems, *Phys. Rev. X* **13**, 021005 (2023).
- [20] Sarah A. M. Loos, Samuel Monter, Felix Ginot, and Clemens Bechinger, Universal symmetry of optimal control at the microscale, *Phys. Rev. X* **14**, 021032 (2024).
- [21] Sergio Ciliberto, Experiments in stochastic thermodynamics: Short history and perspectives, *Phys. Rev. X* **7**, 021051 (2017).
- [22] Juan M. R. Parrondo, Jordan M. Horowitz, and Takahiro Sagawa, Thermodynamics of information, *Nat. Phys.* **11**, 131 (2015).
- [23] Yonggun Jun, Momčilo Gavrilov, and John Bechhoefer, High-precision test of Landauer's principle in a feedback trap, *Phys. Rev. Lett.* **113**, 190601 (2014).
- [24] Antoine Bérut, Artak Arakelyan, Artyom Petrosyan, Sergio Ciliberto, Raoul Dillenschneider, and Eric Lutz, Experimental verification of Landauer's principle linking information and thermodynamics, *Nature (London)* **483**, 187 (2012).
- [25] Marco Ribezzi-Crivellari and Felix Ritort, Large work extraction and the Landauer limit in a continuous Maxwell demon, *Nat. Phys.* **15**, 660 (2019).
- [26] David Abreu and Udo Seifert, Extracting work from a single heat bath through feedback, *Europhys. Lett.* **94**, 10001 (2011).
- [27] Tushar K. Saha, Joseph N. E. Lucero, Jannik Ehrich, David A. Sivak, and John Bechhoefer, Maximizing power and velocity of an information engine, *Proc. Natl. Acad. Sci. U.S.A.* **118**, e2023356118 (2021).
- [28] Tushar K. Saha, Jannik Ehrich, Momčilo Gavrilov, Susanne Still, David A. Sivak, and John Bechhoefer, Information engine in a nonequilibrium bath, *Phys. Rev. Lett.* **131**, 057101 (2023).
- [29] Luca Cocconi and Letian Chen, Efficiency of an autonomous, dynamic information engine operating on a single active particle, *Phys. Rev. E* **110**, 014602 (2024).
- [30] Sudeesh Krishnamurthy, Subho Ghosh, Dipankar Chatterji, Rajesh Ganapathy, and A. K. Sood, A micrometre-sized heat engine operating between bacterial reservoirs, *Nat. Phys.* **12**, 1134 (2016).
- [31] Viktor Holubec, Stefano Steffenoni, Gianmaria Falasco, and Klaus Kroy, Active Brownian heat engines, *Phys. Rev. Res.* **2**, 043262 (2020).
- [32] Étienne Fodor and Michael E. Cates, Active engines: Thermodynamics moves forward, *Europhys. Lett.* **134**, 10003 (2021).
- [33] Thomas Speck, Efficiency of isothermal active matter engines: Strong driving beats weak driving, *Phys. Rev. E* **105**, L012601 (2022).
- [34] Ruben Zakine, Alexandre Solon, Todd Gingrich, and Frédéric Van Wijland, Stochastic stirring engine operating in contact with active baths, *Entropy* **19**, 193 (2017).
- [35] M. J. Schnitzer, S. M. Block, H. C. Berg, and E. M. Purcell, Strategies for chemotaxis, in *Biology of the Chemotactic Response*, edited by J. P. Armitage and J. M. Lackie (Cambridge University Press, Cambridge, 1990), pp. 15–34.
- [36] Mark J. Schnitzer, Theory of continuum random walks and application to chemotaxis, *Phys. Rev. E* **48**, 2553 (1993).
- [37] Michael E. Cates, Diffusive transport without detailed balance in motile bacteria: Does microbiology need statistical physics?, *Rep. Prog. Phys.* **75**, 042601 (2012).
- [38] Rosalba Garcia-Millan and Gunnar Pruessner, Run-and-tumble motion in a harmonic potential: Field theory and entropy production, *J. Stat. Mech.* (2021) P063203.
- [39] Jae Sung Lee and Hyunggyu Park, Effects of the non-Markovianity and non-Gaussianity of active environmental noises on engine performance, *Phys. Rev. E* **105**, 024130 (2022).
- [40] Christopher Jarzynski, Nonequilibrium equality for free energy differences, *Phys. Rev. Lett.* **78**, 2690 (1997).

- [41] Gavin E. Crooks, Nonequilibrium measurements of free energy differences for microscopically reversible Markovian systems, *J. Stat. Phys.* **90**, 1481 (1998).
- [42] Pierre Gaspard and Raymond Kapral, Communication: Mechanochemical fluctuation theorem and thermodynamics of self-phoretic motors, *J. Chem. Phys.* **147**, 211101 (2017).
- [43] Patrick Pietzonka and Udo Seifert, Entropy production of active particles and for particles in active baths, *J. Phys. A* **51**, 01LT01 (2017).
- [44] Thomas Speck, Active Brownian particles driven by constant affinity, *Europhys. Lett.* **123**, 20007 (2018).
- [45] Janik Schüttler, Rosalba Garcia-Millan, Michael E. Cates, and Sarah A. M. Loos, companion paper, Active particles in moving traps: Minimum work protocols and information efficiency of work extraction, *Phys. Rev. E* **112**, 024119 (2025).
- [46] The average potential energy can only be extracted partially, as $(k/2)\langle x(0) \rangle_{v_0}^2 \leq (k/2)\langle x(0)^2 \rangle_{v_0} = \langle V \rangle_{v_0}$, by the Cauchy-Schwarz inequality.
- [47] The diffusive delocalization of the trap center can be prevented by incorporating a recovery step of duration t_{rec} that resets the system to the origin after any N cycles, yielding a process that is cyclic on the level of average quantities with period length $Nt_f + t_{\text{rec}}$. If t_{rec} is chosen to scale sublinearly with N , the corrections to both efficiency and power also scale sublinearly. This implies that all our results derived for the operation without recovery steps remain asymptotically valid for large N . For further details see Ref. [45].
- [48] We can assume that the particle has effectively reached a steady state by the end of the iteration if its duration t_f is greater than the intrinsic relaxation time scales of v and x , τ and k^{-1} , so that we can apply Eq. (9).
- [49] Takahiro Sagawa and Masahito Ueda, Generalized Jarzynski equality under nonequilibrium feedback control, *Phys. Rev. Lett.* **104**, 090602 (2010).
- [50] Solomon Kullback, John C. Keegel, and Joseph H. Kullback, *Topics in Statistical Information Theory* (Springer Science & Business Media, Berlin Heidelberg, 2013), Vol. 42.
- [51] Francisco J. Cao and M. Feito, Thermodynamics of feedback controlled systems, *Phys. Rev. E* **79**, 041118 (2009).
- [52] Takahiro Sagawa and Masahito Ueda, Second law of thermodynamics with discrete quantum feedback control, *Phys. Rev. Lett.* **100**, 080403 (2008).
- [53] Jordan M. Horowitz and Suriyanarayanan Vaikuntanathan, Nonequilibrium detailed fluctuation theorem for repeated discrete feedback, *Phys. Rev. E* **82**, 061120 (2010).
- [54] Yury Polyanskiy and Yihong Wu, Lecture notes on information theory, Lecture Notes for ECE563 (UIUC) (2014), Vol. 6.
- [55] Robert M. Gray, *Entropy and Information Theory* (Springer Science & Business Media, New York, 2011).
- [56] Taking into account the underlying thermodynamic cost of the self-propulsion as a background dissipation term results in a maximum efficiency at some finite frequency $1/t_f$. For example, assuming a constant rate $\mathcal{A}$ of ATP depletion of the self-propulsion mechanism, the efficiency scales like $\eta \propto -\mathcal{W}^{\text{eng}}/(D_v \mathcal{I} + \mathcal{A}t_f)$.
- [57] Arthur Ashkin and James M. Dziedzic, Optical trapping and manipulation of viruses and bacteria, *Science* **235**, 1517 (1987).
- [58] Ivo Buttinoni, Lorenzo Caprini, Laura Alvarez, Fabian Jan Schwarzendahl, and Hartmut Löwen, Active colloids in harmonic optical potentials, *Europhys. Lett.* **140**, 27001 (2022).
- [59] Rémi Goerlich, Luís Barbosa Pires, Giovanni Manfredi, Paul-Antoine Hervieux, and Cyriaque Genet, Harvesting information to control nonequilibrium states of active matter, *Phys. Rev. E* **106**, 054617 (2022).
- [60] <https://github.com/garciamillan/active-information-engine>