

## Semantic Information in a Model of Resource Gathering Agents

Damian R. Sowinski <sup>1,\*</sup>, Jonathan Carroll-Nellenback <sup>1</sup>, Robert N. Markwick <sup>1</sup>, Jordi Piñero <sup>2</sup>, Marcelo Gleiser <sup>3</sup>,  
Artemy Kolchinsky <sup>2,4</sup>, Gourab Ghoshal <sup>1,†</sup> and Adam Frank <sup>1,‡</sup>

<sup>1</sup>Department of Physics and Astronomy, University of Rochester, Rochester, New York 14627, USA

<sup>2</sup>ICREA-Complex Systems Lab, Pompeu Fabra University, 08003 Barcelona, Spain

<sup>3</sup>Department of Physics and Astronomy, Dartmouth College, Hanover, New Hampshire 03755, USA

<sup>4</sup>Universal Biology Institute, The University of Tokyo, 7-3-1 Hongo, Bunkyo-Ku, Tokyo 113-0033, Japan



(Received 12 April 2023; accepted 11 September 2023; published 17 October 2023)

We explore the application of a theory of semantic information to the well-motivated problem of resource foraging. Semantic information is defined as the subset of correlations, which is here measured via the transfer entropy, between agent  $A$  and environment  $E$  that is necessary for the agent to maintain its viability  $V$ . Viability, in turn, is endogenously defined as opposed to the use of exogenous quantities like utility functions. In our model, the forager's movements are determined by its ability to measure, via a sensor, the presence of an individual unit of resource, while the viability function is its expected lifetime. Through “interventions”—scrambling the correlations between agent and environment by noising the sensor—we demonstrate the presence of a critical value of the noise parameter,  $\eta_c$ , above which the forager's expected lifetime is dramatically reduced. On the other hand, for  $\eta < \eta_c$  there is little to no effect on its ability to survive. We refer to this boundary as the *semantic threshold*, quantifying the subset of agent-environment correlations that the agent actually needs to maintain its desired state of staying alive. Each bit of information affects the agent's ability to persist both above and below the semantic threshold. Modeling the viability curve and its semantic threshold via forager and/or environment parameters, we show how the correlations are instantiated. Our work demonstrates the successful application of semantic information to a well-known agent-based model of biological and ecological interest. Additionally, we demonstrate that the concept of semantic thresholds may prove useful for understanding the role information plays in allowing systems to become autonomous agents.

DOI: [10.1103/PRXLife.1.023003](https://doi.org/10.1103/PRXLife.1.023003)

### I. INTRODUCTION

Questions about the role of information in the physics of life extend as far back as Schrödinger's seminal 1944 work, *What Is Life* [1]. Four years later, Shannon published his pioneering work on information theory [2], shortly followed by the discovery that DNA serves as a code for living organisms [3]. These developments have led to a deep interest in the relationship between information, physics, and biology [4]. Since then, the applications of information theory to biology have grown exponentially [5–15], allowing researchers to unpack the ways organisms store and process data about the environment and their own internal states [16–18].

One difficulty with applications of Shannon's information theory to biological systems is its “syntactic” nature. That is, the kinds of extant measures employed in information theory capture statistical correlations between systems without any consideration of the relevance or meaning of those

correlations. Living systems, however, act as agents for whom information is intrinsically meaningful in the most basic sense, that is, whether it can be useful for its *self-production* and *self-maintenance* [19,20]. Life, as a driven, nonlinear, and far-from-equilibrium system, is always in a precarious position and must gather information about the state of the environment and its internal state to endure [21]. Some of this information will be useful for this purpose and some will be irrelevant. In this setting, relevance and meaning can be considered synonymous.

Unlike the well-developed field of syntactic information theory, there exists no widely accepted or applied formal theory of semantic information (previous attempts at developing such a theory include Refs. [22–28]). A goal of a precise mathematical theory of semantic information would be to provide an operational definition useful for characterizing nonlinear far-from-equilibrium systems which can be identified as agents (e.g., organisms or robots). Recently, Kolchinsky and Wolpert [29] (henceforth KW18) developed an explicit formalism for semantic information based on the use of counterfactuals and a notion of *viability*. Their formulation uses the state spaces and probability distributions for an agent  $A$  and its environment  $E$  to characterize the mutual information between the two, while the persistence of  $A$  (its ability to maintain a desired state) is measured through a viability function  $V$ . The concept of meaning here is thus taken in the most basic sense of being related to an

\*Damian.Sowinski@rochester.edu

†gghoshal@pas.rochester.edu

‡afrank@pas.rochester.edu

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.

agent's continued existence. By running *intervened* versions of the system dynamics in which some fraction of the mutual information between agent and environment is scrambled, a formal working definition of the semantic information was characterized in terms of the response of the viability function to such interventions. Importantly, the viability is determined by the inherent coupled dynamics of the system and the environment [29,30], rather than through exogenous utility, cost, error, or loss functions (as is sometimes done when studying the value of information in statistics or in engineering applications [15,31–34]).

The *fitness value of information* has been recently proposed in an ecological setting, investigating the role of information in species growth using population distributions as proxy for probability measures [35]. We note that while the KW18 paper outlines the formalism for semantic information, there is yet to be any application of their framework in a biological setting, or indeed in any realistic individual organism model. As we will see one reason for this is the need to deal with both the formal and computational issues associated with the extensive state spaces associated with the method.

There are several classes of models that describe specific attributes of living systems such as synchronization [36,37], pattern formation [38,39], competition between species for resources [40], stability of ecosystems [41], and simple models of metabolism [42], as well as synthetic cells [43]. Each of these models involves settings where one can test the theoretical framework of semantic information. Indeed, the only work to date that has used the KW18 formalism for semantic information was in the context of synthetic cells [44]. However, in that short study no extension of the formalism was attempted.

In this paper, we make an attempt to test the formalism of semantic information in a model of central interest to biology and ecology: the two-dimensional forager. Our approach is threefold: First, we apply the semantic information framework to the well-motivated problem of resource foraging agents. Second, we develop a detailed mathematical and numerical implementation, operationalizing the original KW18 formalism to explore what simplifying approximations lead to a clear connection between semantic information and the viability of the foraging agents. Finally, we demonstrate the efficacy of our approach in uncovering new insights on the general features of agent-environment dynamics.

Section II focuses on a class of models that jointly address the interrelation between exploration and resource consumption. In such *forager models* [45–56], an agent navigates its environment (using various exploration strategies) in search of resources (food) that it then consumes to maintain its desired internal state (staying alive). The resources are either at fixed locations along the spatial extent of the environment or are replenished at a constant rate in random locations. These dynamics are described in detail in Sec. II A. In such a setting, a natural choice for a viability function is the agent's lifetime, while the environment is simply the field of resources. The correlations between the agent and its environment, which are contingent on the agent's sensorial capability, carry meaningful information about the location of the resources. These correlations are then scrambled by tuning the fidelity of the sensor and then measuring its effect on the viability function, the quantitative details of which make up Sec. II B.

We describe our simulations and demonstrate numerically in Sec. III the existence of a plateau in the viability, capturing the subset of correlations that has little effect on the agent's ability to stay alive. Below this threshold the lifetime of the agent monotonically decays with increased scrambling of the mutual information. Using geometric arguments, we derive an analytical expression for the threshold, relating it directly to sensor resolution. We introduce a semantic information concept, the *viability per bit*, that captures the degree to which each bit of the agent's information on its environment is relevant. This quantity, not articulated in the original KW18 formalism, peaks at the boundary separating the plateau from the decaying region indicating the existence of a *semantic threshold*. Information above the threshold has little to no effect on the agent's viability, while below it, each bit becomes crucial for the agent to stay alive.

On the plateau, we investigate the appearance of a viability lower bound for a wide range of model parameters, then derive its analytical expression. Analytical and numerical evidence supporting these results is agnostic to the particular choice of viability function or indeed the type of foraging strategy. We end in Sec. IV with a discussion of the implications of our findings and possible future directions. Five Appendixes clarify technical aspects of our model and the approximations used.

## II. FORAGING IN A REPLENISHING ENVIRONMENT

Consider a system  $A \times E$ , defined to be a foraging agent,  $A$ , exploring an environment,  $E$ . The states of both agent and environment are described by finite collections of dynamical degrees of freedom,  $a = (a_1, a_2, \dots) \in A$  and  $e = (e_1, e_2, \dots) \in E$ . Evolution in a given time interval is described by a trajectory  $\{(a(t), e(t)) | t \in [t_i, t_f]\}$ . A model is the dynamical rule, deterministic or stochastic, that constructs this trajectory; the subset of the model describing the agent is referred to as a *foraging strategy*.

There are several foraging strategies that have been considered in the literature, including diffusive [45–48], intermittent [49–51], Lévy-flight [52,53], and ballistic trajectories [54–56]. In Sec. II A we define in detail a model of a replenishing environment containing an agent with a ballistic foraging strategy. The choice of this particular strategy stems from its analytical tractability, enabling the derivation of the semantic threshold and viability bound of Sec. III, which are then verified numerically. Other foraging strategies are then considered, and shown numerically to have similar results.

### A. Definition of the model

We begin with a description of our foraging agent, and the environment it navigates. The agent has a fixed metabolic rate  $\mu$  powered by an ever-dwindling fuel reserve, and a sensor allowing it to detect resources within some finite circular range  $R$ . Moving at constant speed  $v$ , the agent changes its direction  $\hat{n}$  only in response to detected resources. Any resource that falls within the collection radius  $r \leq R$  of the agent can be harvested, its energy content added to the fuel supply of the agent. The state of the agent is described by the tuple  $a = (s, \tau, \hat{n}, \mathbf{x})$ .  $s \in [0, S]$  is the agent's stored fuel supply;

TABLE I. Model parameters. M, L, and T are generalized units of mass, length, and time, respectively.

Parameter	Symbol	Units
Collection radius	$r$	L
Detection radius	$R$	L
Speed	$v$	$L T^{-1}$
Metabolic rate	$\mu$	$M L^2 T^{-3}$
Maximum stored energy	$S$	$M L^2 T^{-2}$
Resource energy	$\epsilon$	$M L^2 T^{-2}$
Resource decay rate	$\gamma$	$T^{-1}$
Energy influx	$\Gamma$	$M T^{-3}$

when  $s = 0$  the agent is no longer functional; otherwise it is considered *alive*. The parameter  $\tau \in \{0, 1\}$  is Boolean and describes if the agent has locked onto a target. The agent will not change its direction of motion—indicated by the unit vector  $\hat{\mathbf{n}}$ —if it has a target. Consuming a resource resets  $\tau \mapsto 0$ . Finally,  $\mathbf{x} \in [0, L]^2$  is the position of the agent in the environment, defined as a square of side length  $L \gg R$ .

The state of the environment is specified by the set of locations of a fluctuating number of resources,  $e = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N\}$  with  $\mathbf{y}_n \in [0, L]^2$ ,  $\forall n$ . The available resources are renewable—there is a source of energy flux per unit area into the environment,  $\Gamma$ , resulting in the growth of fixed energy  $\epsilon$  resources across it, which, in turn, decay at a rate  $\gamma$ . In the absence of an agent, the two processes lead to an equilibrium average resource density  $n_{\text{eq}} = \Gamma/\epsilon\gamma$  with  $\mathcal{O}(L^{-1}\sqrt{\Gamma/\epsilon\gamma})$  fluctuations, so that the average distance between resources is  $\ell_{Re} = \sqrt{\epsilon\gamma/\Gamma}$  (see Appendix A for details of the derivation).

Large inhomogeneities in the resource density, such as those which are created by a foraging agent, are repopulated at a rate  $\gamma$ . We work in the small back-reaction regime where forager changes to the equilibrium spacing are on the order  $\delta\ell_{Re}/\ell_{Re} \sim \mathcal{O}(R^2/L^2) \ll 1$ . The agent is then effectively uncorrelated with the environmental degrees of freedom, on distances larger than some  $\mathcal{O}(1)$  multiple of  $R$ . Table I summarizes the parameters that define the model.

As the agent moves through the environment, its collection diameter  $2r$  acts as a cross section, sweeping out an area for harvest per unit time proportional to its speed. Interesting regimes of the model happen when the speed falls within the following bounds:

$$v_* < v < v^* \quad \text{where} \quad v_* = \frac{\epsilon \mu \gamma}{S 2r\Gamma} \quad \text{and} \quad v^* = \frac{\mu \gamma}{2r\Gamma}. \quad (1)$$

For  $v < v_*$ , resources are spread so thin that the foraging area on a full tank has an expected harvest less than 1, leading to agents with rather short lives. For  $v > v^*$ , the harvest area fueled by a single resource has an expected harvest greater than 1: the agent is coupled to an environment of abundance, and is effectively immortal. In between, the agent's probability of survival decays exponentially at long lifetimes, while still forming significant correlations with the environment. This is the regime interesting for exploration. In a strong back-reaction setting these bounds are dynamic: agent harvesting increases the average spacing between resources, thereby increasing both  $v_*$  and  $v^*$ . Again, we will work in the small

back-reaction limit, where the forager does not affect the environment. Figure 1 provides a schematic of the agent-centric dynamics.

## B. The viability function and interventions

While there are a number of ways of quantifying the viability of the agent, an obvious choice is the expected lifetime  $V = \mathbb{E}[T]$ ; here  $\mathbb{E}$  is an expectation value over an ensemble of agents, and  $T$  is the time at which the agent first reaches the dead state ( $s = 0$ ) (as we will later show, our results are robust to other variants of the function). This is easily measured by running an ensemble of agents as described in Fig. 1, and characterizing the distribution of lifetimes. The viability of the agent depends crucially on it correctly setting its trajectory  $\hat{\mathbf{n}}$  when a resource is targeted, which amounts to having a properly working sensor. A broken sensor results in an agent moving in directions that do not correlate with the location of a resource, increasing the chance of starvation and thus death. Importantly, when presenting our results below, we report viability values in units of  $S/\mu$  (maximum stored energy divided by metabolic rate; see Table I). In these units, a viability value of 1 represents the expected lifetime of an agent that starts with a full tank and does not acquire any resources.

To escape the vagaries of chance, the agent needs a sensor that correlates target direction with resource location. Let  $\rho_{AE}(a, e)$  be the joint distribution over states of the agent and environment. We use a standard notation where the subscript identifies the state space that the arguments are drawn from. Without loss of generality, we assume that resources are indexed in order of distance to agent, with  $\mathbf{y}_1$  being the closest. To hone in on the specific correlation between sensor and environment, we next marginalize over all but the closest resource,  $\mathbf{y} = \mathbf{y}_1$ , and condition on the agent being alive, leading to the simpler joint distribution  $\rho_{AE}(\hat{\mathbf{n}}, \tau, \mathbf{y} | s > 0)$ . Rather than computing the correlation between all the agent and environmental degrees of freedom, we now have the simpler task of finding the correlations between the subsets  $\{\hat{\mathbf{n}}, \tau\}$  and  $\{\mathbf{y}\}$ .

Denote the living agent distribution  $\rho_A = \rho_A(\hat{\mathbf{n}}, \tau)$  and the environment distribution  $\rho_E = \rho_E(\mathbf{y})$ . Correlations between agent and environment prevent a factorization of the joint distribution so that  $\rho_{AE} \neq \rho_A \rho_E$ . (The distribution factorizes in the case of a nonfunctional agent, given that the lack of an interaction pathway rapidly suppresses correlations between the agent and environment.) While alive, interactions transfer information between the two, generating the correlations which are then responsible for prolonging the agent's viability. Consequently, we examine the flow of information between the environment and the agent during a detection event.

Consider a resource detection and locking event during the time interval  $\Delta t$ , when a resource is detected at a distance  $r < |\mathbf{y}| \leq R$ . The agent's orientation evolves from  $\hat{\mathbf{n}} \mapsto \hat{\mathbf{n}}'$  while its target degree of freedom flips from  $\tau = 0 \rightarrow \tau' = 1$ . For an agent with a perfect sensor, the transition probability between nontargeting and targeting agent states,  $A \rightarrow A'$ , conditioned on the environment state,  $e$ , is

$$\rho_{A \rightarrow A' | e}(\hat{\mathbf{n}}', \tau' | \hat{\mathbf{n}}, 0, \mathbf{y}) = \delta_{\tau'}^1 \delta(\theta(\hat{\mathbf{n}}', \mathbf{y})), \quad (2)$$

where  $\theta(\mathbf{a}, \mathbf{b})$  is the angle between vectors  $\mathbf{a}$  and  $\mathbf{b}$ ,  $\delta_{\tau'}^1$  is the Kronecker delta, and  $\delta(\theta)$  is the Dirac delta function. By the

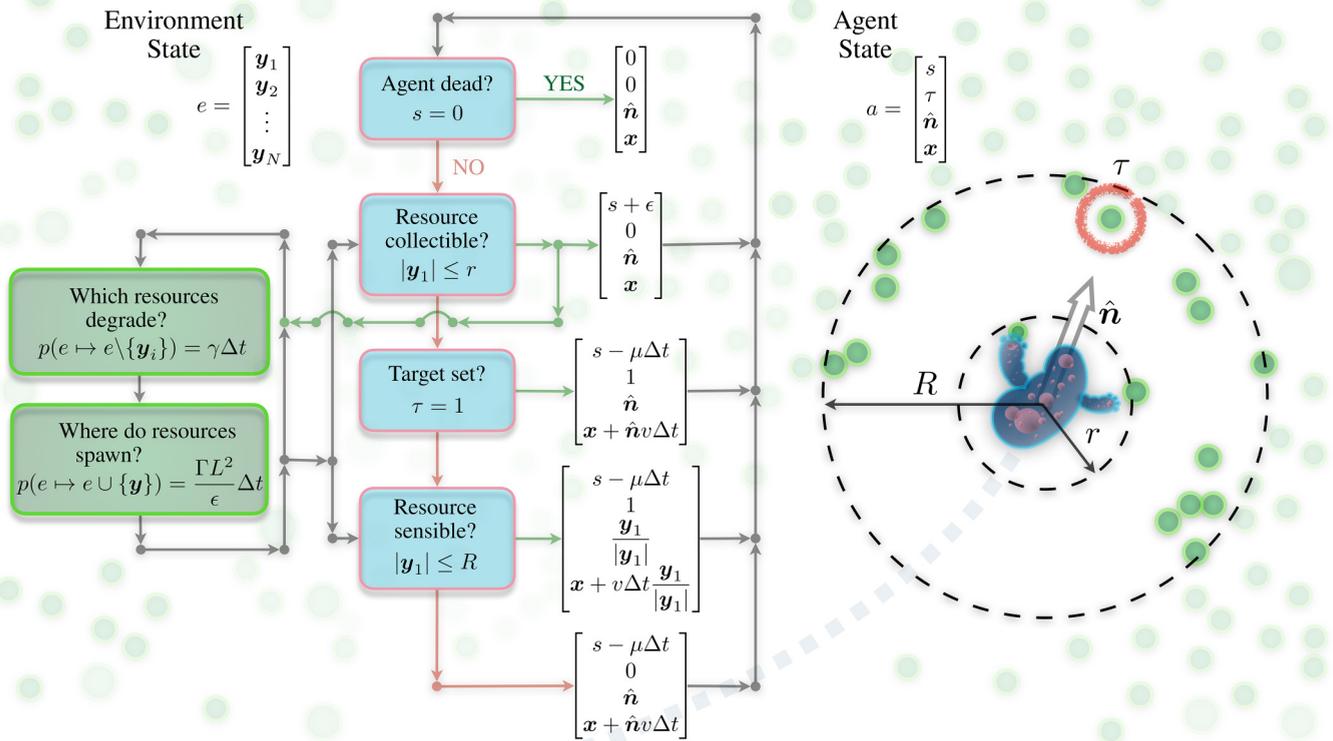


FIG. 1. Overview of model with ballistic strategy. Agent moves at constant speed in a direction  $\hat{n}$ , sensing at a distance  $R$  and collecting resources within radius  $r$ . The two-dimensional environment consists of decaying resources, randomly generated. Agent movement is powered by a metabolism that drains the agent's energy reserves; collecting resources replenishes them. When a resource comes within sensor range, the agent targets (pink circle  $\tau$ ), orients towards it, and then moves until the resource falls within collection range and is consumed.

rotational symmetry of the environment, we also have that  $\rho_{A \rightarrow A'}(\hat{n}', \tau' | \hat{n}, 0) = \delta_{\tau'}^1 (2\pi)^{-1}$ , which is not equal to Eq. (2), indicating that  $\rho_{AE} \neq \rho_A \rho_E$ .

To investigate the extent these correlations are necessary to ensure the agent's longevity, we compromise the agent's ability to find resources by adding noise to its sensor. This will impact the agent's viability function  $V$ . We thus add a small amount of noise  $\eta > 0$  to the agent's ideal sensor, which affects the transition probability for detection defined in Eq. (2) as  $\rho_{A \rightarrow A'}^\eta \equiv \rho_{A \rightarrow A'}^{\eta=0} \mapsto \rho_{A \rightarrow A'}^\eta$ . We measure the effect on the agent's viability as  $\eta$  is increased to 1, where  $\rho_{A \rightarrow A'}^{\eta=1} \equiv \rho_{A \rightarrow A'}$ . Adding noise changes the Dirac delta in Eq. (2) into a uniform distribution of width  $2\pi\eta$ ,  $\delta \mapsto \delta_\eta$ , where

$$\delta_\eta(\theta) = \begin{cases} \frac{1}{2\pi\eta}, & \theta \in [-\pi\eta, \pi\eta] \\ 0 & \text{otherwise.} \end{cases}$$

As the noise parameter  $\eta \rightarrow 0$ , the perfect sensor is recovered, while  $\eta \rightarrow 1$  gives the uniform distribution.

We approximate the information flowing from the environment to the agent during the detection event with the conditional transfer entropy (from the closest resource to the agent's direction, given the agent's state) [57],

$$\mathcal{T}_{E \rightarrow A}^\eta = \mathbb{E}_{\rho_{AE}} \left[ \log_2 \frac{\rho_{A \rightarrow A'}^\eta}{\rho_{A \rightarrow A'}} \right] = \log_2 \frac{1}{\eta}. \quad (3)$$

The information gathered during detection approaches 0 as  $\eta \rightarrow 1$ , preventing the formation of correlations between

agent and environment. It diverges as  $\eta \rightarrow 0$ , which reflects the infinite precision needed to specify a direction in space, as discussed in [58].

### III. RESULTS

Returning to the notion of viability, we extend our definition to the noised sensor channel by defining  $V = V_0 \mapsto V_\eta = \mathbb{E}_\eta[T]$ , where the expectation is now taken over an ensemble of agents with noisy sensors. We interpret the addition of noise as a counterfactual intervention that scrambles the information transfer between an agent with a perfect sensor and its environment. We extract  $V_\eta$  by simulating a single agent within an equilibrated environment at half of full health and measuring how long it survives. This is done  $10^4$  times for each of 200 values of  $\eta \in [0, 1]$ , the first 100 equally spaced in  $[0, 0.15]$  and the second hundred logarithmically spaced in  $[0.15, 1]$ . Further details of the simulation procedure, including the range of values used for the parameters, are described in Appendix B. For animations of several instances of agents with varying values of  $\eta$ , see the Supplemental Material [59].

#### A. Viability and distribution of lifetimes

In the upper panel of Fig. 2, we plot the distribution of lifetimes for three different values of the noise parameter  $\eta$ . The green curve corresponds to the case of a perfect sensor (no scrambling,  $\eta = 0$ ) and sets the baseline distribution of lifetimes which is peaked near the expected lifetime (vertical

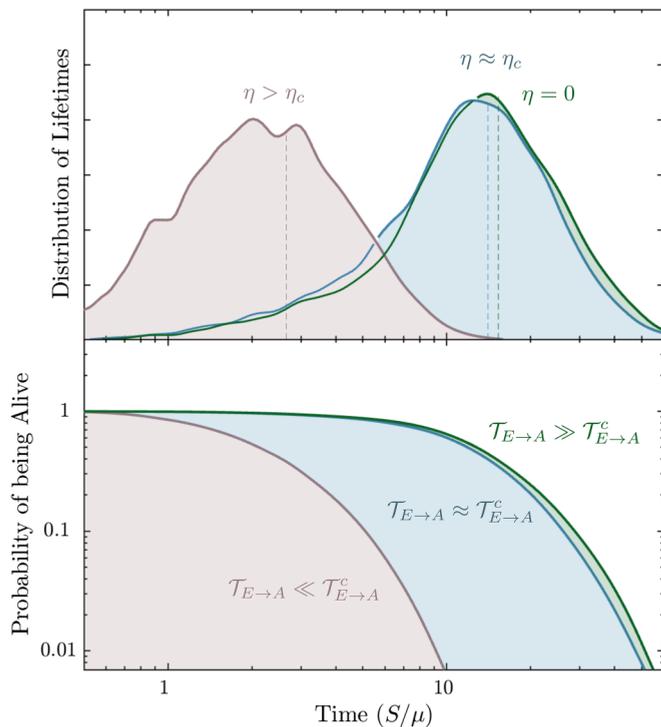


FIG. 2. Top: The distribution of forager lifetimes extracted from the simulations. The green curve is for a noiseless (ideal) agent sensor ( $\eta = 0$ ). The nearly indistinguishable blue curve is from near critical scrambling ( $\eta \approx \eta_c$ ), while the red curve is for above critical scrambling ( $\eta > \eta_c$ ). Bottom: The corresponding distributions for the probability of being alive. Note the decay of the tails is faster than any power law.

dashed dark green line) with a broadly decaying tail. The red curve corresponds to a larger value of  $\eta$  whereby one sees a much lower value for  $\mathbb{E}_\eta[T]$  once a critical value of noise,  $\eta_c$  (blue curve), is surpassed. (Later in the section we derive this threshold analytically.) In the bottom panel, we show the corresponding distributions for the probability of an agent to stay alive, where one can see exponentially decaying tails.

The actual viability  $V_{\text{actual}}$  is defined as the expected lifetime of an agent with noiseless sensor. As anticipated in [29], the viability plateaus near the actual viability  $V_{\text{actual}}$  even as noise degrades the information flow during detection. However, once the sensor efficiency has been sufficiently degraded, the viability begins to drop dramatically, as seen in the top panel of Fig. 3. This rapid decline in viability is interpreted as a semantic threshold—it is the minimal information acquired in a sensing event responsible for maintaining agents close to their actual viability. To better characterize this critical information, we introduce the differential measure of viability per bit (VpB =  $dV/dT$ ), as seen in the lower panel of Fig. 3. The peak in VpB illustrates what an agent requires from its interactions with the environment. Very noisy sensors provide little VpB, even with increased accuracy. However, once a certain level of accuracy is achieved, the VpB grows dramatically and the sensor provides the agent with meaningful information about its environment, resulting in the agent's improved survivability. Too much accuracy, however, is wasted, and the VpB drops again once the sensor goes above and beyond the

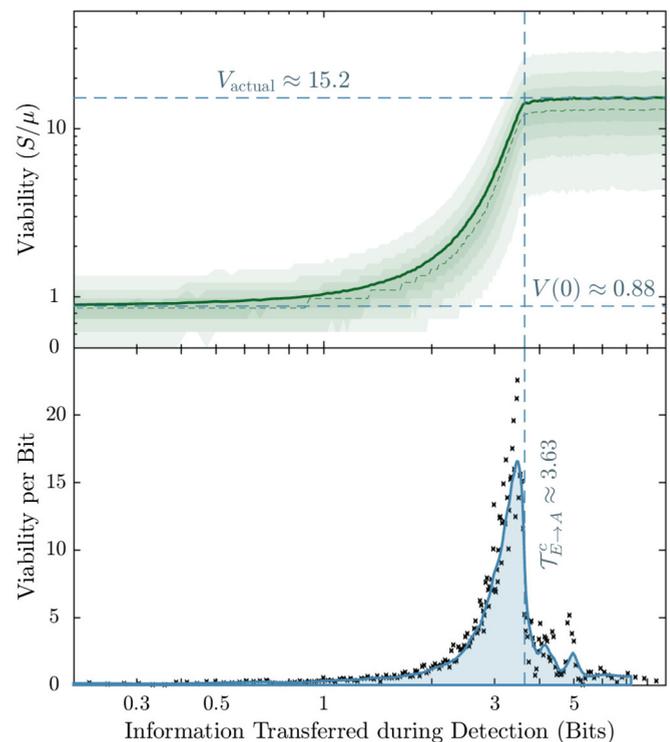


FIG. 3. Top: The viability curve (thick green line) as a function of the transfer entropy,  $\mathcal{T}_{E \rightarrow A}^\eta$ . The actual viability  $V_{\text{actual}}$  is the expected lifetime with no scrambling. The light green line is the median, with the shaded regions representing intervals of 10% above and below the median.  $V(0)$  indicates minimal viability achieved by an agent with a completely noisy sensor, see Eq. (11). Bottom: The viability per bit. The dots are actual data, and the line is a smoothed and interpolated curve added to better visualize the semantic region. The vertical dashed line is the semantic threshold in Eq. (5).

constraints dictated by the physical nature of the agent within the environment.

### B. Robustness to parameters and the viability bound

The results in Fig. 3 do not depend on using the expected lifetime as a viability function; any percentile of the lifetime distribution will do. The dashed green line shows the median of the distribution, while the shaded green regions represent jumps by  $\pm 10\%$  from the median. In all cases the qualitative shape of the curve and the semantic threshold is retained. This feature can be understood by considering how scrambling the sensor relates to the foraging efficacy of the agent. Recall that the collection radius  $r$  acts as an impact parameter, and that targets are set when they enter within the sensing radius  $R$ . For this reason, there is a critical noise parameter above which the agent will sometimes miss its target, and therefore achieve a lower viability. This critical value, and the corresponding semantic threshold, are

$$\eta_c = \frac{1}{\pi} \sin^{-1} \frac{r}{R}, \quad (4)$$

$$\mathcal{T}_{E \rightarrow A}^c = \log_2 \pi - \log_2 \sin^{-1} \frac{r}{R}. \quad (5)$$

(See Appendix C for details.) Interestingly, this expression does not depend on most parameters that characterize the agent, such as its speed and metabolic rate, only on its geometry. There is low VpB when the information transferred during a detection event is well above  $\mathcal{T}_{E \rightarrow A}^c$ , and high VpB as  $\mathcal{T}_{E \rightarrow A}^n \rightarrow \mathcal{T}_{E \rightarrow A}^c$ .

Above the semantic threshold ( $\eta < \eta_c$ ), the minute drop in the viability plateau away from  $V_{\text{actual}}$  also has a geometric origin. Due to the circular shape of both the resource collection and sensing zones, a sensor with  $\eta < \eta_c$  will force agents to travel farther ever so slightly to collect a targeted resource. Assume an agent targets a resource a distance  $y = |\mathbf{y}| \leq R$  away, but the sensor is mistaken by an angle  $\theta$  as to the direction of the resource. Rather than having to travel  $y - r$  to collect its target, the agent now has to travel  $y \cos \theta - \sqrt{r^2 - y^2 \sin^2 \theta}$ . Given an information transfer of  $\mathcal{T}$  during detection, averaging over all angles we see that resources a distance  $y$  from the agent are effectively a distance  $(1 + \lambda)y$  with

$$\lambda = \frac{r}{y} \left( 1 - \frac{2^{\mathcal{T}}}{\pi} E \left( \frac{\pi}{2} \middle| \frac{y^2}{r^2} \right) \right) - \left( 1 - \frac{2^{\mathcal{T}}}{\pi} \sin \frac{\pi}{2^{\mathcal{T}}} \right). \quad (6)$$

Here  $E(\varphi|k^2)$  is the indefinite elliptic integral of the second kind with modulus  $k = y/r$ . (See Appendix D for details of the derivation.) We may consider two interesting limiting regimes. In the regime of accurate sensors,  $\eta \approx 0$  and  $\mathcal{T} \gg 1$ , the dilation decays exponentially as

$$\lambda \sim \frac{\pi^2}{6} \frac{y - r}{r} 4^{-\mathcal{T}}. \quad (7)$$

Another interesting regime to consider is when  $y$  is very close to (but above) the collection radius  $r$ . In this case,

$$\lambda \sim \frac{y - r}{r} \left( \frac{\tanh^{-1} \tan \left( \frac{\pi}{2} 2^{-\mathcal{T}} \right)}{\frac{\pi}{2} 2^{-\mathcal{T}}} - 1 \right), \quad (8)$$

up to first order in  $y - r$ . As expected, the dilation grows linearly away from  $r$ , with the slope being controlled by the transfer entropy.

The expression for the dilation factor  $\lambda$  indicates that, above the semantic threshold ( $\eta < \eta_c$ ), agents can be thought of as unscrambled agents making their way through a distorted world. As the information acquisition from the environment is restored ( $\mathcal{T} \rightarrow \infty$ ), the distortion decays exponentially as expected. Of course, within the extraction radius there is no dilation as the scrambled agent still has the reach  $y = r$ , even if it cannot judge properly distances  $y > r$ . Maximal dilation,  $\lambda^{\text{max}}$ , occurs at  $y = R$  in Eq. (6). This constitutes an upper bound, which enables the determination of a lower bound on the viability of a scrambled agent, given the actual viability of an unscrambled agent.

To compare the viability of scrambled and unscrambled agents, we scale down the length scale of the scrambled agent by  $(1 + \lambda^{\text{max}})$  so that it matches the length scale of the unscrambled agent. Unfortunately, this means the speeds of the two no longer match; to restore equality we must scale down the timescale by the same amount. Numerical simulations indicate that this rescaling does not affect the environmental variables. Focusing therefore only on the agent variables, the

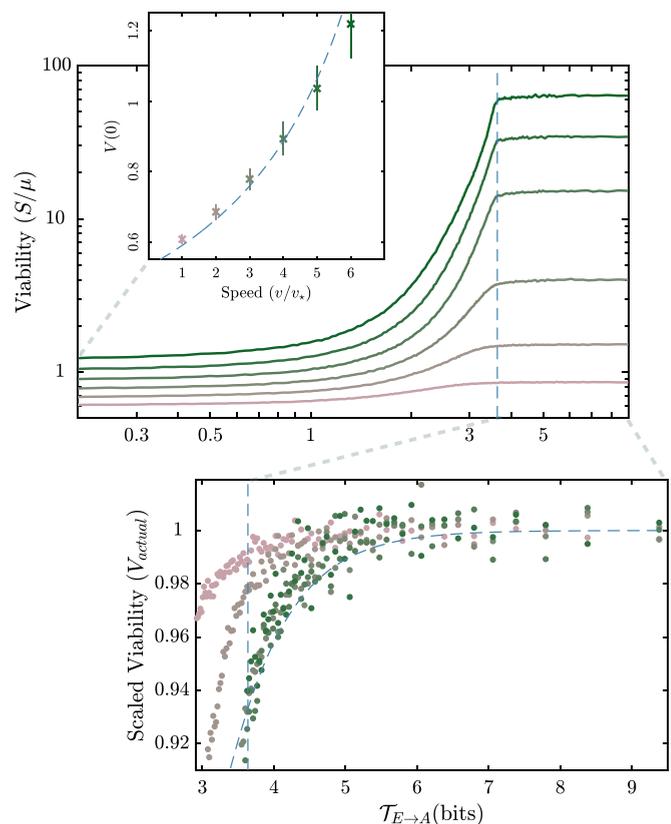


FIG. 4. Top: Viability curves for six velocities in the range  $1 \leq v/v_* \leq 6$  in integer increments. Here  $v_* = 10v_c$ . The vertical dashed line is Eq. (5). Inset: Minimal viability as achieved by an agent with a completely noisy sensor, as  $\mathcal{T}_{E \rightarrow A} \rightarrow 0$ . Bars represent  $5\sigma$  errors in the estimates. The dashed line represents Eq. (11), showing good agreement with simulations. Bottom: The plateau region of each curve, rescaled for purposes of comparison. The dashed curve represents the viability bound in Eq. (10).

rescaled viability is

$$\begin{aligned} V_\eta &= \mathbb{E}_\eta[T] = \mathbb{E}_0[(1 + \lambda^{\text{max}})^{-1} T] \\ &= (1 + \lambda^{\text{max}})^{-1} V_{\text{actual}}. \end{aligned} \quad (9)$$

Finally, expanding this expression for  $\lambda^{\text{max}} \ll 1$ , above the semantic threshold, the viability as a function of the transfer entropy is bounded from below by

$$\begin{aligned} \frac{V(\mathcal{T})}{V_{\text{actual}}} &\gtrsim 2 - \frac{r}{R} + \frac{2^{\mathcal{T} - \mathcal{T}^c}}{\sin^{-1} \frac{r}{R}} \\ &\times \left[ \frac{r}{R} E \left( \frac{\sin^{-1} \frac{r}{R}}{2^{\mathcal{T} - \mathcal{T}^c}} \middle| \frac{R^2}{r^2} \right) - \sin \left( \frac{\sin^{-1} \frac{r}{R}}{2^{\mathcal{T} - \mathcal{T}^c}} \right) \right], \end{aligned} \quad (10)$$

where  $\mathcal{T}^c$  is defined in Eq. (5).

It is instructive to check whether the results presented thus far are robust to varying the speed of the forager. As shown in Eq. (1), the speed is naturally bounded in the limits  $v_* < v < v^*$ . In our simulations, we set  $v^* = 10v_c$ ; varying the velocity within this range, in the upper panel of Fig. 4, we show the viability curves for six values of  $1 \leq v/v_* \leq 6$  in integer increments (colors represent different values of the ratio). The vertical dashed line is the analytically calculated semantic

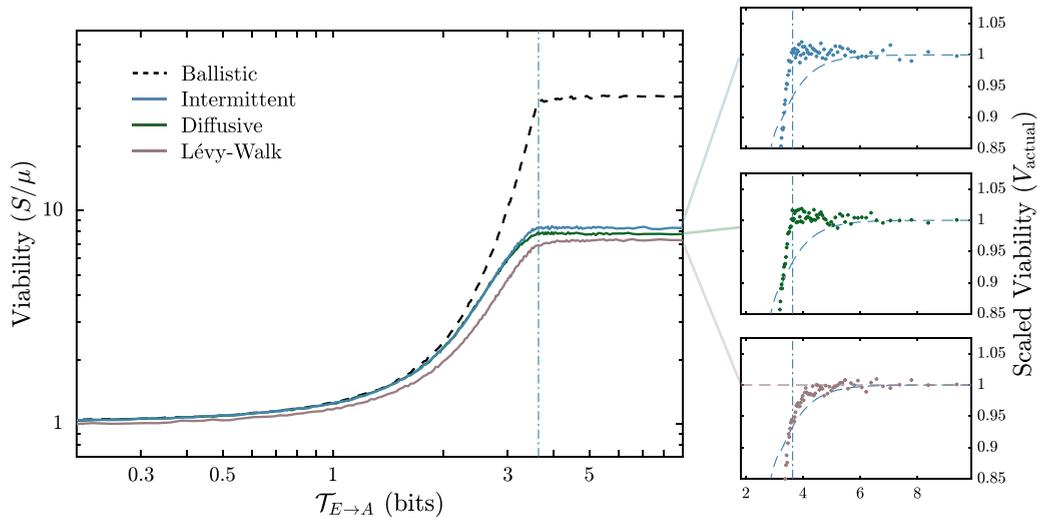


FIG. 5. Comparison of the semantic threshold and viability bound for multiple foraging strategies. Left: The (dashed) viability curve for the ballistic forager is presented for reference. In all strategies foragers move with a velocity  $5v_*$  (in the Lévy flight  $\langle v \rangle = 5v_*$ ). Right: Rescaled viability data points in the plateau region. The vertical (dot-dashed) line is the semantic threshold from Eq. (5), and the dashed curve is the viability bound from Eq. (11).

threshold [Eq. (5)]. As expected, while the height of the viability curve decreases with decreasing velocity (eventually flattening near the lower bound  $v_*$ ), the semantic threshold as well as the shape of each curve is robust to the change in velocities.

In the bottom panel we zoom in to the plateau region for each curve, rescaling them for purposes of comparison. The viability bound calculated in Eq. (10) is plotted as a dashed curve, while the semantic threshold is shown as a vertical dashed line. While the agreement with the viability bound improves with increasing velocity, Eq. (10) bounds the plateau for all velocity ranges (up to statistical fluctuations).

One can estimate the minimal viability,  $V(0) = \mathbb{E}_1[T]$ , as achieved by an agent with a completely noisy sensor. To do so, note that resources encountered by such an agent follow a Poisson distribution with rate  $2vr/\ell_{Re}^2$ . The total fuel processed by the agent will therefore be  $s_0 + \epsilon \mathbb{E}_1[T]$ , where  $s_0$  is the initial stored fuel. But this total fuel divided by the metabolic rate is precisely the expected lifetime; a little algebra reveals

$$V(0) = \frac{s_0/\mu}{1 - v/v^*}, \tag{11}$$

where  $v^*$  is the upper velocity limit defined in Eq. (1). As mentioned earlier, the agent approaches immortality when its velocity approaches this limit, even if it has a completely random sensor. In the inset panel of Fig. 4, we plot Eq. (11) as a dashed line. Points indicate the minimal viability extracted from simulations as a function of velocity, showing excellent agreement with the theoretical prediction.

### C. Robustness to choice of foraging strategy

To see how robust these results are to changes in the foraging behavior, we consider several other strategies of increasing complexity. The most oft-studied version is that of a diffusive forager that undergoes a random walk when not detecting a

resource (and then moving ballistically on detection) [45–48]. Adding an additional layer of complexity are intermittent foragers [49–51]. Here, the forager moves diffusively, but not always for the purpose of detecting resources. For large periods the movement represents locomotion, with intermittent sensing and consumption events. This represents the fact that biological organisms are not always looking to consume, but spend their time performing other functions. Finally, Lévy-walk foragers [52,53] are inspired by empirical observations of animal movement [60], whereby foragers move at variable speed drawn from a (truncated) power-law distribution of velocities. In Appendix E, we provide details on how each of these strategies is incorporated into our simulations.

The viability curve for each strategy is plotted in the left-hand panel of Fig. 5, where the dashed line indicates ballistic, the blue intermittent, the green diffusive, and finally the pink line Lévy-walk foraging. We note that the shape of the curve for all strategies is retained, and that the analytically calculated semantic threshold matches the numerical simulations for all considered strategies. In the right-hand panel we show the (rescaled) plateau region for all strategies indicating good agreement with the viability bound, Eq. (11). The figure indicates that the height of the viability plateau decreases with increasingly complex strategies, with ballistic foraging leading to the highest viability and Lévy walk the lowest. This is a pleasing result, given the previous observation that ballistic foraging is the most *efficient* in a replenishing environment [54] (where efficiency is measured in terms of the encounter rate of the forager with a resource).

Thus, the viability factorizes into a plateau that depends implicitly on both environment and agent, and a scaling term that depends only on properties of the agent. Furthermore, below the semantic threshold, the viability curves for multiple strategies asymptote to the same minimum viability. This is fairly straightforward to understand. When locked onto a target, every implementation of our strategies defaults to ballistic motion, and maintains it until a resource is collected.

For the range of simulation parameters used, there is a high probability of a resource being found in the sensing area of a newly instantiated agent ( $P_{N>0} \approx 99.2\%$ ) so it is most likely that the instance has  $\tau = 1$  and is moving ballistically. The different strategies change what foragers do when not locked onto a target. The formation of low-resource “deserts” is therefore crucial to switching between behaviors that differentiate strategies, but as sensor scrambling is increased, agents tend to travel farther each time they engage in ballistic motion, and end up in areas well outside their local desert. Then, once again, there is a high probability that a resource is within sensing range, and the process begins anew. In the limit of a completely noisy sensor, all variants behave like a ballistic forager that chooses a random direction after each resource collection.

#### IV. DISCUSSION AND CONCLUSION

In this work, we explored the application of a theory of semantic information to the well-motivated problem of a resource foraging agent. Semantic information is defined as the subset of correlations—measured here via the transfer entropy  $\mathcal{T}_{E \rightarrow A}^{\eta}$ —between agent  $A$  and environment  $E$  that is necessary for the agent to maintain its viability  $V$ . Viability, in turn, is endogenously defined as opposed to the use of exogenous quantities like utility functions (we note a subtlety in terms of different possible definitions of “system existence,” which could be a topic of further exploration). The semantic information content in a particular agent-environment system is determined by rerunning the system evolution for intervened versions of the dynamics. Such interventions involve “scrambling” the information transfer between the agent and environment in a specified way. Tracking changes in agent viability for a distribution of intervened system trajectories allows the semantic information content to be determined.

Applying this procedure to our forager model required finding appropriate approximations to the fully specified state space of the original KW18 formalism. This “operationalization” of the KW18 prescription in terms of demonstrating how to apply it to a multidimensional model is, itself, an important result of our study. For realistic systems with many degrees of freedom, the full specification of the joint probability distribution  $\rho_{AE}$  over the full state space, as required in the initial formalism, may prove computationally expensive or even intractable. For our model we adopted a phenomenological perspective which traced over those degrees of freedom not essential to specification of the viability (this is, in part, what we mean by “operationalization”). This left us with a reduced state space in which intervened trajectories could be simulated and the subsequent transfer entropy could be calculated. This represents the application of semantic information in a (quasi)realistic setting.

In our model, the forager’s movements were determined by its ability to sense the presence of an individual unit of resource. Once detected (sensing was limited to within a radius  $R$ ), the forager moved towards the resource. The transfer entropy was scrambled by adding noise to the forager’s sensor via a parameter  $\eta$ , where  $\eta = 1$  implied a complete loss of the forager’s ability to sense the direction of the resource. Our results, expressed in terms of a viability function defined as

the expectation value of the forager’s lifetime, clearly showed the effect of adding noise to the sensor. For  $\eta > \eta_c$ , where  $\eta_c$  is a critical value set by the sensing radius, the forager’s expected lifetime was dramatically reduced (Fig. 2). We refer to this as the semantic threshold.

This result by itself represents an important extension of previous work on forager dynamics [45–52,56]. What is different in our work is the way in which casting the problem in terms of semantic information reveals useful aspects of the model dynamics. For instance, while ballistic foraging has been shown to be the most efficient strategy in terms of optimizing the encounter rate of foragers with resources (in a replenishing environment) [54], here we find the same result but recast in terms of information theory and the concept of viability (or survival). The transfer entropy represents correlations established between the forager (an agent) and the environment via the agent’s sensor. A blind forager  $\eta = 1$  has a sensor that is fully decoupled from the environment. By tracking how the forager’s viability changes as these correlations are either increased or decreased, we gain some understanding about the role they play in the forager’s ability to persist. In particular the upper panel of Fig. 3, which shows the viability curve  $V(\mathcal{T})$ , reveals two essential ways to understand the role of such correlations for the forager (or for any agent) which we now unpack separately.

Above the semantic threshold (i.e., to the right of threshold in Figs. 3–5) we find a plateau of high viability. Moving from right to left in this region we are removing correlations between agent and environment; however, this does not effect the agent’s ability to maintain its existence. Thus the correlations that are being removed are not essential to maintain agent viability. Below the semantic threshold, each bit of information affects the agent’s ability to persist. Once this threshold is passed, we see the viability monotonically decrease. Thus casting the forager-environment system into the semantic information formalism allows us to see exactly how much information matters. In addition, our ability to model the shape of the  $V(\mathcal{T})$  curve (including in particular the location of the semantic threshold) in terms of the forager and environment parameters  $(R, r, \mu, \eta)$  allows us to see how the correlations are instantiated. Layering on increasingly complex foraging strategies mimicking various aspects of biological systems leads to the same qualitative results. The semantic threshold, the shape of the viability curve, and the viability bound are robust to whether the forager moves ballistically, diffusively, intermittently, or by undertaking a Lévy flight. Thus, our work can provide a useful starting point for studies of more complex biological behavior in terms of semantic information in order to better understand the underlying structure of correlations between agents and their environments as well as the role information dynamics plays in their behavior.

Coming from the left in Fig. 3 leads us to a different perspective, which may prove useful in using semantic information to understand how agents arise in the first place. Beginning with the low-viability region on the left, we see that adding correlations initially has little effect. The forager dies quickly and adding an additional bit of correlation with the environment does not change that outcome. As more bits of information are acquired, the agent’s viability slowly rises.

However, it is only near the semantic threshold that the slope of the curve accelerates and *viability per bit* (VpB) peaks. Thus, it is possible that this threshold can prove useful for understanding the role information plays in allowing systems to become autonomous agents. Both hurricanes and cells are nonlinear, driven, far-from-equilibrium systems, but only cells are considered agents. Future work could explore the relationship between the accumulation of semantic information and the emergence of agentlike behavior while also considering the thermodynamic cost of such accumulation.

### ACKNOWLEDGMENTS

This project was partly made possible through the support of Grant No. 62417 from the John Templeton Foundation. The opinions expressed in this publication are those of the author(s) and do not necessarily reflect the views of the John Templeton Foundation. The authors thank the Center for Integrated Research Computing (CIRC) at the University of Rochester for providing computational resources and technical support. A.K. thanks Sosuke Ito for support and encouragement. J.P. is supported by “María de Maeztú” fellowship MDM-2014-0370-17-2.

### APPENDIX A: EQUILIBRIUM RESOURCE DENSITY

Consider a large patch of the environment,  $A \subset \mathbb{R}^2$ . The number of resources in the patch is

$$N(t) = \int_A d^2x n(\mathbf{x}, t), \quad (\text{A1})$$

where  $n$  is the resource density. Resources do not move, but they can be created or destroyed. Creation is due to an energy flux impinging on  $A$ , denoted  $\Gamma(\mathbf{x}, t)$ ; there exists some mechanism in the environment that converts this flux into localized resource deposits of energy  $\epsilon$ . Destruction is due to another mechanism by which resources decay, their energy lost as heat, denoted  $\gamma(\mathbf{x}, t)$ . Note that the former rate is independent of the number of resources; the latter is not. The rate of change of resources in this large patch is

$$\frac{dN}{dt} = \int_A d^2x \left( \frac{\Gamma(\mathbf{x}, t)}{\epsilon} - \gamma(\mathbf{x}, t)n(\mathbf{x}, t) \right). \quad (\text{A2})$$

Taking  $\Gamma$  and  $\gamma$  both to be static, the resource density satisfies

$$\begin{aligned} \frac{\partial n}{\partial t} &= \frac{\Gamma}{\epsilon} - \gamma n \\ \Downarrow \\ n(t) &= n_0 e^{-\gamma t} + n_{eq}(1 - e^{-\gamma t}), \end{aligned} \quad (\text{A3})$$

where the equilibrium density is  $n_{eq} = \Gamma/\epsilon\gamma$ . The average area occupied by a single resource is the reciprocal of this,  $\epsilon\gamma/\Gamma$ , so that the average spacing between resources at equilibrium is

$$\ell_{Re} = \sqrt{\frac{\epsilon\gamma}{\Gamma}}. \quad (\text{A4})$$

If instead we coarse-grain over the positions of resources and care only about their count, then the transition probabilities for a resource being generated or destroyed in a time

TABLE II. Primary simulation dimensionless parameters.

	Parameter	Symbol	Value
Scaling	Collection radius	$\tilde{r}$	1
	Metabolic rate	$\tilde{\mu}$	1
	Maximum stored energy	$\tilde{S}$	1
	Detection radius	$\tilde{R}$	4
	Speed	$\tilde{v}$	20 (Fig. 3), 25 (Fig. 5)
Chosen	Time discretization	$\Delta\tilde{t}$	0.02
	Environment size	$\tilde{L}$	100
	Resource energy	$\tilde{\epsilon}$	0.1
	Resource decay rate	$\tilde{\gamma}$	0.1
	Energy influx	$\tilde{\Gamma}$	0.001
Derived	Minimal speed	$\tilde{v}_*$	5
	Maximal speed	$\tilde{v}^*$	50
	Equilibrium resource density	$\tilde{n}_{eq}$	0.1
	Equilibrium resource spacing	$\tilde{\ell}_{Re}$	$\sqrt{10} \approx 3.16$

$\Delta t \rightarrow 0$  are  $p(N \rightarrow N+1) = \Gamma A \Delta t / \epsilon$  and  $p(N \rightarrow N-1) = \gamma N \Delta t$ . The transitions between these coarse-grained states define an  $M/M/\infty$  queue—a well-known stochastic process with a stationary distribution that is Poisson:

$$p(N) = \frac{1}{N!} \left( \frac{\Gamma A}{\epsilon \gamma} \right)^N e^{-\frac{\Gamma A}{\epsilon \gamma}}. \quad (\text{A5})$$

The expected number of resources is  $\langle N \rangle = \frac{\Gamma A}{\epsilon \gamma}$ , which, under the assumption of a homogeneous distribution, gives an equilibrium number density that matches  $n_{eq}$  from the preceding paragraphs. The variance in the number of resources is  $\delta N^2 = n_{eq} A$ , giving a fluctuation in the number density of  $\delta n = \ell_{Re}^{-1} A^{-1/2}$ . Thus when considering large environments,  $A \gg \ell_{Re}^2$ , the relative fluctuations around the equilibrium density are negligible,  $\delta n/n \rightarrow 0$ .

### APPENDIX B: SIMULATION DETAILS

Translating our model for numerical simulation requires first dimensionalizing all parameters in Table I. Natural scales were chosen to be agent-centric, so that numerical values within each simulation are all interpreted in terms of scales meaningful to the forager. The natural timescale is  $T = S/\mu$ , the lifetime of an agent on a full tank—with no other influxes of fuel; the natural length scale is  $L = r$ , the extraction radius of the agent; and the natural mass scale,  $M = S T^2 L^{-2}$ , the size of a full fuel tank.

Dimensionalization is accomplished by dividing model parameters by the combination of natural units given in the far right column of Table I. For example, the dimensionless (tilded) speed is  $\tilde{v} = v/(L T^{-1})$ , read off naturally as the number of extraction radii ( $\sim$  forager length) traveled on a full tank of fuel. We summarize the chosen scales and dimensionless values of parameters, as well as derived dimensionless values of relevant quantities, in Table II.

After dimensionalizing, data are generated by simulating  $10^4$  instances of a single forager placed into an environment at equilibrium, for multiple values of the scrambling parameter

$\eta$ . Specifically, 200 values of  $\eta$  were chosen, with the first 100 spaced linearly on the interval  $[0, 0.15)$ , and the second 100 logarithmically spaced on the interval  $[0.15, 1]$ . These were chosen after initial coarser simulations revealed the location of the viability plateau. For each value of  $\eta$ , each of the  $10^4$  simulations are begun by instantiating a forager at half health,  $\tilde{s} = 0.5$ , at the center of an environment—a square with side length  $\tilde{L} = 100$ —with  $N = \tilde{n}_{eq}\tilde{L}^2 = 10^3$  resources distributed uniformly.

Each individual simulation evolves with a temporal discretization  $\Delta\tilde{t} = 1/50 = 0.02$  following the flowchart presented in Fig. 1. Below we summarize the chart, describing what our code does in a single time step for both the degrees of freedom (DOF) of the agent,  $(s, \tau, \hat{n}, \mathbf{x})$ , and environment,  $(y_1, y_2, \dots)$  (whose DOF are ordered by distance from the agent). First, the agent must be alive,  $s > 0$ ; otherwise the simulation ends. Otherwise the closest resource,  $y_1$ , is examined. The probability that two or more resources have the same closest distance is for all intents and purposes zero, but should this miracle occur only one among these is chosen, uniformly at random. If that resource is within the extraction radius of the agent, then it is harvested, whereupon the agent's fuel supply increases  $s \mapsto s + \epsilon$ , the agent's target DOF is set to  $\tau \rightarrow 0$ , while its orientation and position remain unchanged,  $(\hat{n}, \mathbf{x}) \mapsto (\hat{n}, \mathbf{x})$ . If that resource is not within extraction range, then the agent's target DOF,  $\tau$ , is examined. If a target is set,  $\tau = 1$ , then it remains set, as does the orientation of the agent; simultaneously the fuel supply decreases,  $s \mapsto s - \mu\Delta t$ , as the agent moves,  $\mathbf{x} \mapsto \mathbf{x} + v\hat{n}\Delta t$ . If the target is not set,  $\tau = 0$ , then the code checks if the closest resource is within sensing range of the agent. If it is, then the agent sets a target  $\tau \mapsto 1$ , and expends fuel  $s \mapsto s - \mu\Delta t$ , to reorient and move towards it,  $(\hat{n}, \mathbf{x}) \mapsto (y_1/|y_1|, \mathbf{x} + y_1/|y_1|\Delta t)$ . If the closest resource is beyond sensing range, then the target remains unset, and the agent's orientation does not change as it expends fuel,  $s \mapsto s - \mu\Delta t$ , to move in that direction,  $\mathbf{x} \mapsto \mathbf{x} + \hat{n}v\Delta t$ .

Since we have a finite environment, sometimes the agent may find itself approaching a boundary, in which case the code implements reflective boundary conditions as follows. Should one of the components of the position of the agent fall outside the interval  $[0, L]$ , then it is replaced by the reflected position in the interval. That is, if  $x < 0$  then  $x \mapsto -x$ , while if  $x > L$  then  $x \mapsto L - x$ . If a spatial reflection takes place, then the corresponding component of the velocity is also reflected,  $v_x \mapsto -v_x$ . This completes the agent's update time step.

Next, the environment updates via a full time step consisting of a degradation step followed by a growth step. During the degradation step, each resource is checked to see if it degrades. A check consists of drawing a uniform random number and seeing if it is greater than  $e^{-\gamma\Delta t}$ . If it is, the resource is removed from the environment; otherwise it remains. During the growth step, a random non-negative integer is chosen from a Poisson distribution with  $\lambda = \Gamma L^2 \Delta t / \epsilon = 2$ , i.e.,

$$p_n = \frac{e^{-\lambda}}{n!} \lambda^n. \quad (\text{B1})$$

This integer represents how many new resources are generated; these are added to the environment at uniformly random

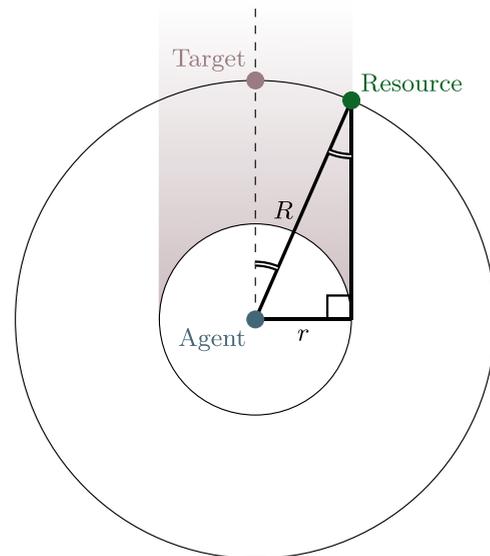


FIG. 6. The geometry of an agent (mis)targeting a resource at the edge of sensor range.

locations. The displacement and distance from the agent to each resource is updated, and the list of resource locations is reordered according to proximity. This completes the environment update time step.

For each of the  $10^4$  simulations at each value of  $\eta$ , the degrees of freedom of the agent and the location of the nearest resource are saved at each time step. On the plateau, simulations run for  $\sim 40$  minutes, averaging  $\sim 750$  time steps and resulting in data sets of  $\sim 0.5$  Gb; simulations on the other side of the plateau dropped in both temporal and spatial complexity. For a fixed set of parameters,  $10^4$  simulations across the 200 scrambling parameters require  $\sim 8000$  core-hours of computation and generate  $\sim 60$  Gb of data. Simulations were conducted on the Center for Integrated Research Computing's BlueHive cluster at the University of Rochester.

### APPENDIX C: THE SEMANTIC THRESHOLD

Consider an agent for which a resource has just entered sensor range and a targeting event has occurred. We scramble the information relayed by the event so that the agent reorients itself with a target that is misaligned with the resource. The agent begins its journey towards the target, as shown in Fig. 6. For an agent moving towards a target any resource that falls within the grey region will be collected. The maximal misalignment angle between target and resource is  $\pi\eta$ . From the diagram one can infer a critical level of scrambling,  $\eta_c$ , satisfying

$$\sin \pi \eta_c = \frac{r}{R}, \quad (\text{C1})$$

below which the agent will always collect the resource. Above it, the agent will sometimes miss the resource—an event that could potentially result in starvation. Combining Eq. (C1) with Eq. (3) gives the semantic threshold, Eq. (5).

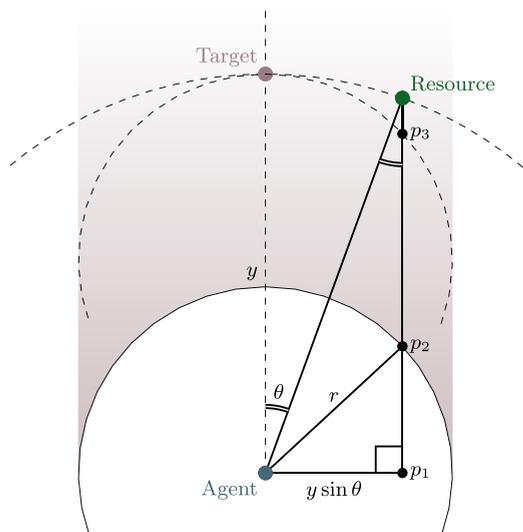


FIG. 7. The geometry describing how far an agent has to travel to collect a mistargeted resource.

#### APPENDIX D: THE VIABILITY PLATEAU

Let us examine a case below the critical scrambling strength, so that every targeted resource is still collected. Once again we consider a scrambled targeting event, but assume it occurs for a resource a distance  $y < R$  away from the agent. This type of event occurs regularly after the agent has collected a resource and there is at least one other resource in sensor range. The agent begins to move towards the target as depicted in Fig. 7.

Had the sensor not been scrambled, the agent would traverse a distance  $y - r$  to collect the resource. Meanwhile, the scrambled agent needs to travel the distance from  $p_2$  to the resource for collection to occur; a little geometry and trigonometry on the diagram give us this distance as  $y \cos \theta - \sqrt{r^2 - y^2 \sin^2 \theta}$ . Since  $p_2$  to  $p_3$  is also a distance  $y - r$ , denote the remaining distance  $\lambda(\theta)y$  where  $\lambda(\theta)$  is referred to as a *dilation* factor. It will become clear why in what follows.

A little algebra gets us the angular dependence of the dilation factor,

$$\lambda(\theta) = \frac{r}{y} - 1 + \cos \theta - \sqrt{\frac{r^2}{y^2} - \sin^2 \theta}. \quad (\text{D1})$$

Of course this factor will be different each time a targeting event occurs, depending on distance to resource and the particular mismatch angle of the target. However, for a given  $\eta$  we know the distribution of mismatch angles—it is uniform over  $[-\pi\eta, \pi\eta]$ —so for fixed  $y$ , the expected dilation is

$$\mathbb{E}_\eta[\lambda] = \frac{1}{2\pi\eta} \int_{-\pi\eta}^{\pi\eta} d\theta \lambda(\theta).$$

Examining the functional form of  $\lambda(\theta)$  we see that the first two terms are trivial and the third term easily integrates to  $\sin \pi\eta/\pi\eta$ , while the fourth term requires a little massaging. After factoring out a  $-r/y\pi\eta$  and using that the integrand is

even, the remaining factor is massaged into the form

$$\int_0^{\pi\eta} d\theta \sqrt{1 - \frac{y^2}{r^2} \sin^2 \theta}.$$

The integral is not elementary, but it turns out to be an incomplete elliptic integral of the second kind,

$$E(\varphi|k^2) = \int_0^\varphi d\theta \sqrt{1 - k^2 \sin^2 \theta}, \quad (\text{D2})$$

with modulus  $k$ . There is a slight nuance since normally  $0 < k^2 < 1$ , while in our case  $y > r$  gives a modulus greater than unity. This bound ensures the integrand remains real for all values of  $\varphi$ . Fortunately, since we are considering cases below critical scrambling,  $y \sin \theta < r$  is automatically satisfied and there is no need to worry about which branch we are on. With this minutiae out of the way, we write the expected dilation for fixed  $y$ ,

$$\mathbb{E}_\eta[\lambda] = \frac{r}{y} - 1 + \frac{\sin \pi\eta}{\pi\eta} - \frac{r}{y} \frac{E(\pi\eta|\frac{y^2}{r^2})}{\pi\eta}, \quad (\text{D3})$$

which, when combined with Eqs. (3) and (5), yields Eq. (6). To simplify the notation, we simply denote this as  $\lambda$  in the main text.

The asymptotic behaviors of the dilation require Taylor expanding the elliptic integral. For the perfect sensor limit we treat  $\eta \ll 1$ , so that

$$\begin{aligned} E(\pi\eta|k^2) &\approx \int_0^{\pi\eta} 1 - \frac{k^2}{2} \theta^2 \\ &= \pi\eta - \frac{k^2}{6} \pi^3 \eta^3 \end{aligned}$$

to third order in  $\eta$ . Doing the same to the sine function, and plugging into Eq. (D3), one finds the first nonvanishing term is second order in  $\eta$ :

$$\frac{\pi^2}{6} \frac{y - r}{r} \eta^2.$$

Plugging in the inverse of Eq. (3) gives the quoted exponential asymptotic behavior.

The asymptotic when  $y \rightarrow r^+$  requires a bit more work. We parametrize the approach as  $y = (1 + \varepsilon)r$  with  $\varepsilon \ll 1$ , and expand the elliptic integral to first order:

$$\begin{aligned} E(\pi\eta|(1 + \varepsilon)^2) &= \int_0^{\pi\eta} d\theta \sqrt{1 - (1 + \varepsilon)^2 \sin^2 \theta} \\ &\approx \int_0^{\pi\eta} d\theta \sqrt{\cos^2 \theta - 2\varepsilon \sin^2 \theta} \\ &\approx \int_0^{\pi\eta} d\theta \cos \theta - \varepsilon \int_0^{\pi\eta} d\theta \cos \theta \tan^2 \theta \\ &= (1 - \varepsilon) \sin \pi\eta - \varepsilon \ln |\sec \theta + \tan \theta|. \end{aligned}$$

After plugging this in, Eq. (D3) simplifies to

$$\varepsilon \left( \frac{1}{\pi\eta} \ln \frac{\cos \pi\eta}{1 + \sin \pi\eta} - 1 \right).$$

The argument of the logarithm can be massaged using elementary trigonometric identities into  $(1 - \tan \frac{\pi\eta}{2})/(1 + \tan \frac{\pi\eta}{2})$ .



$\epsilon$  and higher speeds, a greater distinction appeared between the curves, with the plateau of the intermittent strategy lying between the diffusive and ballistic strategies.

### 3. Lévy-walk forager

The final model goes back to the ballistic forager and elevates the speed parameter to a dynamical degree of freedom. The speed is a stochastic multiple of  $v_0$ , drawn from a power law with exponent  $\nu$ , namely,

$$P(k; k_{\min}, k_{\max}, \nu) = \frac{k^{-\nu}}{\sum_{n=k_{\min}}^{k_{\max}} n^{-\nu}}, \quad (\text{E1})$$

meant to simulate a random Lévy walk. This novelty is added at the end of the ballistic forager time step, updating the velocity to some multiple,  $k_{\min} \leq k \leq k_{\max}$ .

For our experiments, we wanted to compare this strategy to the others in the limit of a completely noisy sensor. Initial experiments revealed that this required the average velocity of all the models to be the same in order for the viabilities to match in that limit (the precise reason is elucidated in the text). To this end we chose  $\nu = 1.9658$ , with  $k_{\min} = 2$ ,  $k_{\max} = 30$ , meaning that the maximal velocity achieved by the agent was six times  $v^*$ . These ensured

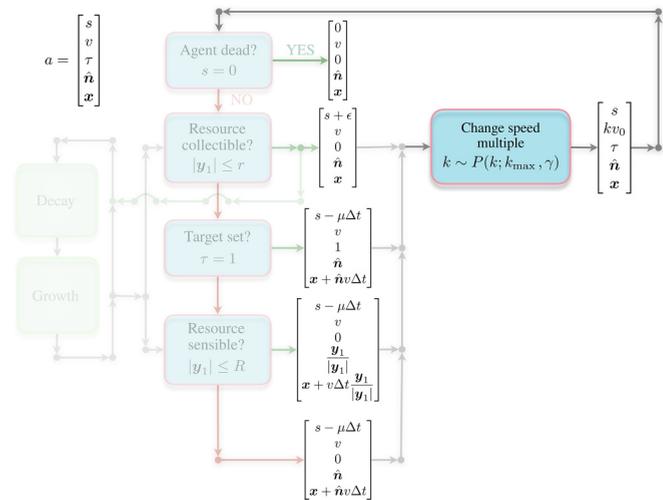


FIG. 10. Lévy-walk foraging strategy.

that  $\langle v \rangle / v_* = 5.0$ , the same as the other simulations, to four decimal places of accuracy. Our implementation is shown in Fig. 10.

- [1] E. Schrodinger, *What Is Life* (Cambridge University Press, Cambridge, UK, 1944)
- [2] C. E. Shannon, A mathematical theory of communication, *Bell Syst. Tech. J.* **27**, 379 (1948).
- [3] J. D. Watson and F. H. Crick, Molecular structure of nucleic acids: A structure for deoxyribose nucleic acid, *Nature (London)* **171**, 737 (1953).
- [4] J. Hopfield, Physics, computation, and why biology looks so different, *J. Theor. Biol.* **171**, 53 (1994).
- [5] P. Godfrey-Smith, Information in biology, in *The Cambridge Companion to the Philosophy of Biology*, Cambridge Companions to Philosophy, edited by D. L. Hull and M. Ruse (Cambridge University Press, Cambridge, UK, 2007), pp. 103–119.
- [6] H. Mattingly, K. Kamino, B. Machta, and T. Emonet, *Escherichia coli* chemotaxis is information limited, *Nat. Phys.* **17**, 1426 (2021).
- [7] B.-O. Küppers, *Information and the Origin of Life* (MIT Press, Cambridge, MA, 1990).
- [8] C. Adami, Information theory in molecular biology, *Phys. Life Rev.* **1**, 3 (2004).
- [9] S. Uda, Application of information theory in systems biology, *Biophys. Rev.* **12**, 377 (2020).
- [10] A. Gohari, M. Mirmohseni, and M. Nasiri-Kenari, Information theory of molecular communication: Directions and challenges, *IEEE Trans. Mol. Biol. Multi-Scale Commun.* **2**, 120 (2016).
- [11] A. Rhee, R. Cheong, and A. Levchenko, The application of information theory to biochemical signaling systems, *Phys. Biol.* **9**, 045011 (2012).
- [12] G. Tkačik and W. Bialek, Information processing in living systems, *Annu. Rev. Condens. Matter Phys.* **7**, 89 (2016).
- [13] M. C. Donaldson-Matasci, C. T. Bergstrom, and M. Lachmann, The fitness value of information, *Oikos* **119**, 219 (2010).
- [14] O. Rivoire and S. Leibler, The value of information for populations in varying environments, *J. Stat. Phys.* **142**, 1124 (2011).
- [15] R. M. Hazen, P. L. Griffin, J. M. Carothers, and J. W. Szostak, Functional information and the emergence of biocomplexity, *Proc. Natl. Acad. Sci. USA* **104**, 8574 (2007).
- [16] D. R. Sowinski, J. Carroll-Nellenback, J. DeSilva, A. Frank, G. Ghoshal, and M. Gleiser, The consensus problem in politics of agents with dissimilar cognitive architectures, *Entropy* **24**, 1378 (2022).
- [17] S. Oh, E. F. Bowen, A. Rodriguez, D. Sowinski, E. Childers, A. Brown, L. Ray, and R. Granger, Towards a perceptual distance metric for auditory stimuli, [arXiv:2011.00088](https://arxiv.org/abs/2011.00088).
- [18] E. F. Bowen, A. M. Rodriguez, D. R. Sowinski, and R. Granger, Visual stream connectivity predicts assessments of image quality, *J. Vision* **22**, 4 (2022).
- [19] G. Schlosser, Self-re-production and functionality, *Synthese* **116**, 303 (1998).
- [20] M. Mossio, C. Saborido, and A. Moreno, An organizational account of biological functions, *Br. J. Philos. Sci.* **60**, 813 (2009).
- [21] X. Fang and J. Wang, Nonequilibrium thermodynamics in cell biology: Extending equilibrium formalism to cover living systems, *Annu. Rev. Biophys.* **49**, 227 (2020).
- [22] D. Polani, T. Martinetz, and J. Kim, An information-theoretic approach for the quantification of relevance, in *Advances in Artificial Life: 6th European Conference, ECAL 2001 Prague, Czech Republic, September 10–14, 2001 Proceedings 6* (Springer, Berlin, 2001), pp. 704–713.
- [23] E. Thompson and M. Stapleton, Making sense of sense-making: Reflections on enactive and extended mind theories, *Topoi* **28**, 23 (2009).
- [24] C. L. Nehaniv, D. Polani, K. Dautenhahn, R. te Beekhorst, and L. Cañamero, Meaningful information, sensor evolution, and the temporal horizon of embodied organisms, in *Proceedings*

- of the Eighth International Conference on Artificial Life (ICAL 2003)* (MIT Press, Cambridge, MA, 2002), pp. 345–349.
- [25] J. Barham, A dynamical model of the meaning of information, *Biosystems* **38**, 235 (1996).
- [26] T. W. Deacon, Shannon-Boltzmann-Darwin: Redefining information (part I), *Cognitive Semiotics* **1**, 123 (2007).
- [27] P. A. Corning, Control information theory: The ‘missing link’ in the science of cybernetics, *Syst. Res. Behav. Sci.* **24**, 297 (2007).
- [28] M. Gleiser and D. Sowinski, How we make sense of the world: Information, map-making, and the scientific narrative, in *The Map and the Territory: Exploring the Foundations of Science, Thought and Reality* (Springer, Berlin, 2018), pp. 141–164.
- [29] A. Kolchinsky and D. H. Wolpert, Semantic information, autonomous agency and non-equilibrium statistical physics, *Int. Focus.* **8**, 20180041 (2018).
- [30] C. Rovelli, Meaning and intentionality = information + evolution, in *Wandering Towards a Goal*, The Frontiers Collection (Springer, Cham, 2018), pp. 17–27.
- [31] J. P. Gould, Risk, stochastic preference, and the value of information, *J. Econ. Theory* **8**, 64 (1974).
- [32] R. L. Stratonovich, *Theory of Information and its Value*, edited by R. V. Belavkin, P. M. Pardalos, and J. C. Principe (Springer, Cham, 2020).
- [33] C. E. Shannon *et al.*, Coding theorems for a discrete source with a fidelity criterion, *IRE Nat. Conv. Rec.* **4**, 1 (1959).
- [34] D. R. Sowinski, M. D. McGarry, E. E. Van Houten, S. Gordon-Wylie, J. B. Weaver, and K. D. Paulsen, Poroelasticity as a model of soft tissue structure: Hydraulic permeability reconstruction for magnetic resonance elastography in silico, *Front. Phys.* **8**, 617582 (2021).
- [35] J. Usinowicz and M. I. O’Connor, The fitness value of ecological information in a variable world, *Ecol. Lett.* **26**, 621 (2023).
- [36] J. A. Acebrón, L. L. Bonilla, C. J. Pérez Vicente, F. Ritort, and R. Spigler, The Kuramoto model: A simple paradigm for synchronization phenomena, *Rev. Mod. Phys.* **77**, 137 (2005).
- [37] D. García-Selfa, G. Ghoshal, C. Bick, J. Pérez-Mercader, and A. P. Muñuzuri, Chemical oscillators synchronized via an active oscillating medium: Dynamics and phase approximation model, *Chaos, Solitons Fractals* **145**, 110809 (2021).
- [38] F. Cooper, G. Ghoshal, A. Pawling, and J. Pérez-Mercader, Internal composite bound states in deterministic reaction diffusion models, *Phys. Rev. Lett.* **111**, 044101 (2013).
- [39] S. Mimar, M. M. Juane, J. Park, A. P. Muñuzuri, and G. Ghoshal, Turing patterns mediated by network topology in homogeneous active systems, *Phys. Rev. E* **99**, 062303 (2019).
- [40] J. E. Cohen, Ratio of prey to predators in community food webs, *Nature (London)* **270**, 165 (1977).
- [41] N. Rooney, K. McCann, G. Gellner, and J. C. Moore, Structural asymmetry and the stability of diverse food webs, *Nature (London)* **442**, 265 (2006).
- [42] J. C. Xavier, W. Hordijk, S. Kauffman, M. Steel, and W. F. Martin, Autocatalytic chemical networks at the origin of metabolism, *Proc. R. Soc. B: Biol. Sci.* **287**, 20192377 (2020).
- [43] J. C. Blain and J. W. Szostak, Progress toward synthetic cells, *Annu. Rev. Biochem.* **83**, 615 (2014).
- [44] B. Ruzzante, L. Del Moro, M. Magarini, and P. Stano, Synthetic cells extract semantic information from their environment, *IEEE Trans. Mol. Biol. Multi-Scale Commun.* **9**, 23 (2023).
- [45] O. Bénichou and S. Redner, Depletion-controlled starvation of a diffusing forager, *Phys. Rev. Lett.* **113**, 238101 (2014).
- [46] U. Bhat, S. Redner, and O. Bénichou, Does greed help a forager survive? *Phys. Rev. E* **95**, 062119 (2017).
- [47] O. Bénichou, M. Chupeau, and S. Redner, Role of depletion on the dynamics of a diffusing forager, *J. Phys. A: Math. Theor.* **49**, 394003 (2016).
- [48] U. Bhat and S. Redner, How smart should a forager be? *J. Stat. Mech.: Theory Exp.* (2022) 033402.
- [49] O. Bénichou, C. Loverdo, M. Moreau, and R. Voituriez, Intermittent search strategies, *Rev. Mod. Phys.* **83**, 81 (2011).
- [50] O. Bénichou, M. Coppey, M. Moreau, P. Suet, and R. Voituriez, A stochastic model for intermittent search strategies, *J. Phys.: Condens. Matter* **17**, S4275 (2005).
- [51] O. Bénichou, C. Loverdo, M. Moreau, and R. Voituriez, Two-dimensional intermittent search processes: An alternative to Lévy flight strategies, *Phys. Rev. E* **74**, 020102(R) (2006).
- [52] F. Bartumeus, J. Catalan, U. L. Fulco, M. L. Lyra, and G. M. Viswanathan, Optimizing the encounter rate in biological interactions: Lévy versus Brownian strategies, *Phys. Rev. Lett.* **88**, 097901 (2002).
- [53] G. Viswanathan, F. Bartumeus, S. V. Buldyrev, J. Catalan, U. Fulco, S. Havlin, M. Da Luz, M. L. Lyra, E. Raposo, and H. E. Stanley, Lévy flight random searches in biological phenomena, *Physica A* **314**, 208 (2002).
- [54] A. James, M. J. Plank, and R. Brown, Optimizing the encounter rate in biological interactions: Ballistic versus Lévy versus Brownian strategies, *Phys. Rev. E* **78**, 051128 (2008).
- [55] D. Campos, J. Cristín, and V. Méndez, Optimal escape-and-feeding dynamics of random walkers: Rethinking the convenience of ballistic strategies, *Phys. Rev. E* **103**, 052109 (2021).
- [56] F. Bartumeus, D. Campos, W. S. Ryu, R. Lloret-Cabot, V. Méndez, and J. Catalan, Foraging success under uncertainty: Search tradeoffs and optimal space use, *Ecol. Lett.* **19**, 1299 (2016).
- [57] T. Schreiber, Measuring information transfer, *Phys. Rev. Lett.* **85**, 461 (2000).
- [58] M. Gleiser and D. Sowinski, Configurational information approach to instantons and false vacuum decay in  $D$ -dimensional spacetime, *Phys. Rev. D* **98**, 056026 (2018).
- [59] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PRXLife.1.023003> for videos of the ballistic forager model below, near, and above the critical scrambling parameter.
- [60] H. Barbosa, M. Barthelemy, G. Ghoshal, C. R. James, M. Lenormand, T. Louail, R. Menezes, J. Ramasco, F. Simini, and M. Tomasini, Human mobility: Models and applications, *Phys. Rep.* **734**, 1 (2018).
- [61] M. Egbert, M. M. Hanczyc, I. Harvey, N. Virgo, E. C. Parke, T. Froese, H. Sayama, A. S. Penn, and S. Bartlett, Behaviour and the origin of organisms, *Origins Life Evol. Biospheres* **53**, 87 (2023).