Sense organ control in moths to moles is a gamble on information through motion

Chen Chen[1,2], **Todd D. Murphey**[1,3], **Malcolm A. MacIver**[1,2,3,4*]

***For correspondence:** maciver@northwestern.edu ( )

[1]The Neuroscience and Robotics Laboratory, Northwestern University, Evanston IL;
[2]Department of Biomedical Engineering, Northwestern University, Evanston IL;
[3]Department of Mechanical Engineering, Northwestern University, Evanston IL;
[4]Department of Neurobiology, Northwestern University, Evanston IL

Abstract Sensory organs—be they independently movable like eyes or requiring whole body movement as in the case of electroreceptors—are actively manipulated throughout stimulus-driven behaviors. While multiple theories for these movements exist, such as infotaxis, in those cases where they are sufficiently detailed to predict sensory organ trajectories, they show poor fit to measured trajectories. Here we present evidence that during tracking, these trajectories are predicted by energy-constrained proportional betting, where the probability of moving a sense organ to a location is proportional to an estimate of how informative that location will be combined with its energetic cost. Energy-constrained proportional betting trajectories show good agreement with measured trajectories of four species engaged in visual, olfactory, and electrosensory tracking tasks. Our approach combines information-theoretic approaches in sensory neuroscience with analyses of the energetics of movement. It can predict sense organ movements in animals and prescribe them in robotic tracking devices.

Introduction

Sensory organs undergo small lateral movements as they near or hold station close to a target of interest (***Martin, 1965***; ***Basil et al., 2000***; ***Ferner and Weissburg, 2005***; ***Webb et al., 2004***; ***Willis and Avondet, 2005***; ***Porter et al., 2007***; ***Louis et al., 2008***; ***Duistermars et al., 2009***; ***Yovel et al., 2010***; ***Khan et al., 2012***; ***Stamper et al., 2012***; ***Catania, 2013***; ***Sponberg et al., 2015***; ***Lockey and Willis, 2015***; ***Rucci and Victor, 2015***; ***Stockl et al., 2017***) (Fig. 1). There are several models in the literature that have been proposed (***Stamper et al., 2012***; ***Yovel et al., 2010***; ***Khan et al., 2012***; ***Rucci and Victor, 2015***; ***Najemnik and Geisler, 2005***; ***Yang et al., 2016***; ***Stockl et al., 2017***). For example, in the related case of signal-emitter organ control, fruit bats are known to oscillate their tongue-click-based sonar signals on approach to their targets (***Yovel et al., 2010***). This can be effective because for many signal sources, the signal intensity peaks at the target's location and tapers away in all directions. The expected amount of information—in the bat study quantified by the Fisher Information of the emitted sonar signal—is highest at the maximum slope of the signal profile because at those locations, small variations in the emitter position leads to large changes in emitted signal power on target and thus also in the returning echos. In contrast, at the flat peak of the profile where the object is located, small variations in emitter position lead to small or no change in signal and returning echo; the expected information is therefore low. For active sensing animals like bats, dolphins, and electric fish, placing emitter organs so that the target is at a location of high

signal slope then leads to better information harvesting and hence better estimation of the target location (***Clarke et al., 2015***; ***Yovel et al., 2010***). The same is true for animals guided by light or sound, through placement of sense organs at high slope locations. Puzzlingly, this would suggest that animals should monitor an information peak (one location of high signal slope), while the documented animal behavior suggests that they move between multiple information peaks (***Yovel et al., 2010***) (one peak for each slope). One approach, called infotaxis (***Vergassola et al., 2007***), similarly generates trajectories that, at each move on a grid of locations, maximizes information (Fig. 2B). As we will show, however, infotactic trajectories show poor fit to what animals do during tracking behavior.

Here we propose that sensory organs (or emitters in the case of active sensing animals) are moved according to a very different principle: rather than move the sensor to maximize information, move it to sample spatially distributed signals proportionate to the expected information density (EID), constrained by the energetic cost of movement. The underlying sensory sampling strategy gambles on the chance of obtaining more information at a given location through carefully controlled sensor motion that balances two factors that typically push in opposite directions: 1) proportionally bet on the expected information gain; and 2) minimize the energy expended for motion.

To understand proportional betting on information, consider a distribution of expected information for an animal tracking an object on the ground, visualized as a color map in Fig. 2A. The illustrated trajectory—as determined by the motor control decisions of the animal—proportionally bets on the EID: the sensor takes more frequent samples (dots along trajectory) in information-rich regions and samples sparsely to provide coverage in information-poor regions. Similarly, in behaviors where animals sample discretely over time, sampling proportional to the EID might take the form of varying the sampling frequency or location at which samples are taken, as observed in bats, beaked whales, humans, and pulse electric fish (***Yovel et al., 2010***; ***Kothari et al., 2018***; ***Caputi et al., 2003***; ***Pluta and Kawasaki, 2008***; ***Nelson and MacIver, 2006***; ***Schnitzler et al., 2003***; ***Madsen et al., 2005***; ***Yang et al., 2016***; ***Hoppe and Rothkopf, 2019***). The energy constraint means that this trajectory will be executed with finite speed, varying from slow in information-rich regions, to fast in information-sparse regions. For the fruit bat behavior discussed above, rather than the bat's oscillatory movement between two information peaks of the sonar beam (their Fig. S5 and Fig. 2A, ***Yovel et al.*** (***2010***)) being an unexplained "bug" of information maximization, it is an expected feature of proportional betting; similar oscillations will be seen throughout the subsequent measured and simulated target tracking behaviors of this study.

For this study we have quantified the expected informativeness of sensing locations by how much an observation at a location would reduce the Shannon-Weaver entropy (hereafter entropy) of the current estimate of the target's location, as in infotaxis (***Vergassola et al., 2007***). (Other measures of information such as Fisher Information can be used with near identical results (***Miller et al., 2016***). In our approach, the closeness of a given trajectory to perfect proportional betting is quantified by the ergodic metric. The ergodic metric provides a way of comparing a trajectory to a distribution (*i.e.*, the EID) by asking whether a trajectory over some time interval has the same spatial statistics as a given distribution (Methods, Appendix 3).Comparing a trajectory to a distribution is a novel capability of the ergodic metric (***Mathew and Mezic, 2011***) not shared by common methods of comparing two probability distributions (Appendix 2). Through extremizing ergodicity we obtain trajectories that bet on information, expending a metabolic cost to move to informative locations in the space. With a perfectly ergodic trajectory (one with an ergodic measure of zero, only possible with infinite time and when the energy of movement is not considered), the distribution of expected information is perfectly encoded by the trajectory, or, equivalently, the trajectory does perfect proportional betting on the EID. We therefore call our approach ergodic information harvesting (hereafter EIH).

Figure 2C–E illustrates the emergence of oscillatory sensing-related motions in one-dimensional tracking behavior simulated with EIH, here for a hypothetical moth tracking a flower swaying lat-

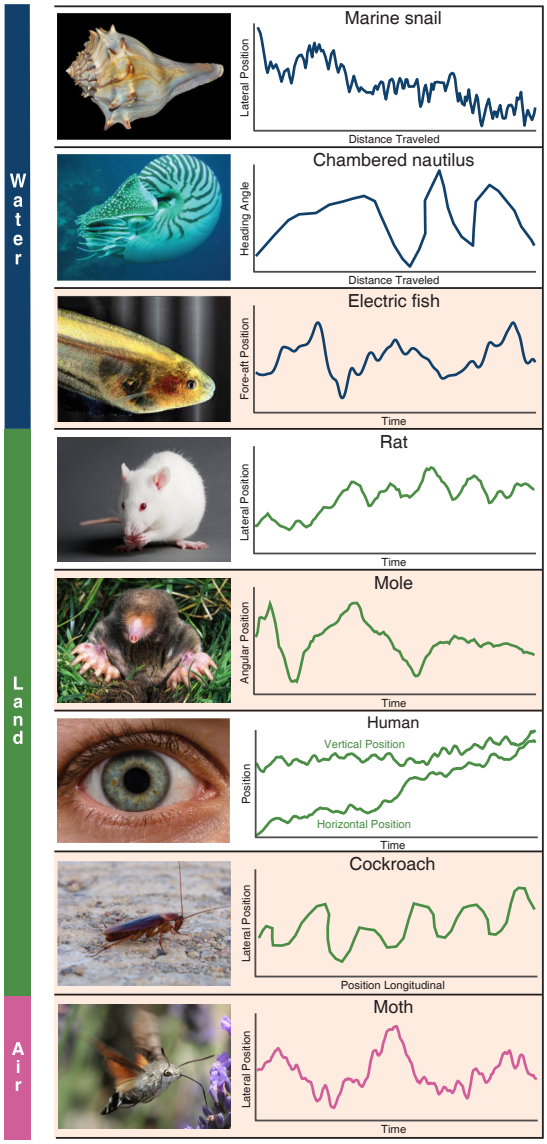

Figure 1. Whole-body or sensory organ small-amplitude motions are ubiquitous as animals track targets. Here we show siphon casting behavior in the marine snail (body) (***Ferner and Weissburg, 2005***), cross-current swimming in the Chambered nautilus (body) (***Basil et al., 2000***), whole-body oscillations in electric fish (body position) (***Stamper et al., 2012***), zigzagging motions in the rat (nose) (***Khan et al., 2012***), back and forth searching in the mole (nose) (***Catania, 2013***), fixational eye movements in humans (eye) (***Rucci and Victor, 2015***), the zigzagging walk of the cockroach (head) (***Willis and Avondet, 2005***), and flower tracking motions in the moth (head) (***Stockl et al., 2017***). These sensing-related movements occur with striking consistency across animals using different sensory modalities and within different physical environments. We propose these movements arise as a form of gambling on information through motion, and show evidence in support of this claim from the four species highlighted (electric fish, mole, cockroach, and moth).

erally in the wind in order to feed from the flower's nectary (***Stockl et al., 2017***). A key behavioral signature of EIH—increased sensor wiggling as the signal weakens—is evident in this illustration. In addition to proportional betting being used at the cognitive decision-making level in primates (***Monosov et al., 2015***; ***Gottlieb et al., 2014***), we suggest that it occurs more broadly as a sensori-motor routine across a wide phylogenetic bracket. Below we will show evidence for this claim by comparing measured tracking trajectories to those simulated with EIH, when EIH is provided the original target trajectory for two species of insect, a fish, and a mammal.

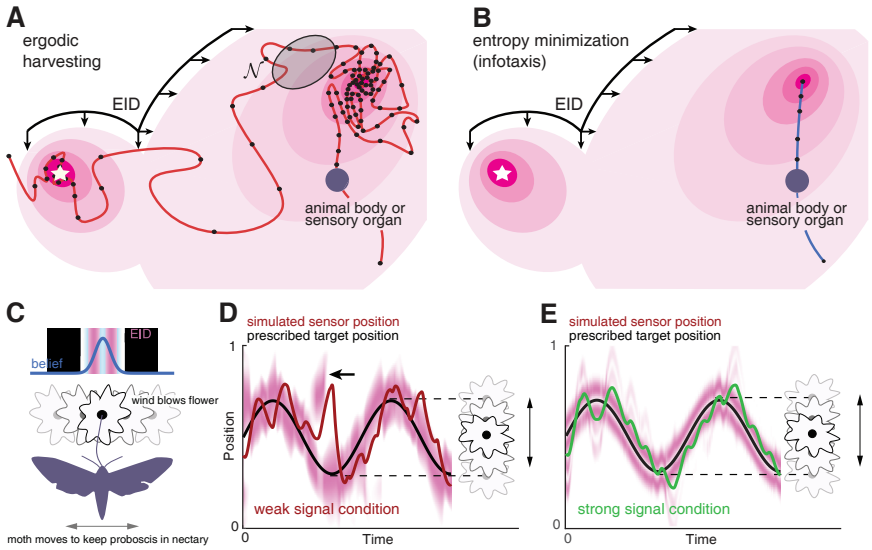

Figure 2. Proportional betting with respect to the cost of gambling on information through motion. (**A–B**) Illustration of an animal localizing a target on the ground by moving its sense organs (*e.g.,* eye, nose, antenna, whole body *etc.*) along a simulated proportional betting trajectory from Miller et al. (***Miller et al., 2016***). (**A**) Trajectory generated with proportional betting on the expected information density (EID, pink–red heatmap). Unknown to the animal, the real target is near the white star on the left which marks the true location of peak information density, while on the right is a distractor (at center of red zone). A distractor in the EID could be caused by either a physically present object sensed by the animal or, in the context of this paper, unmodeled uncertainties such as false-positive sensory input. Each step of our energy-constrained proportional betting algorithm is taken so as to optimally balance two factors, one representing how much that step increases the similarity of the trajectory to one that does perfect proportional betting (where the percentage of time the trajectory spends in any subset $\mathcal{N}$ from $t = 0$ to $t = T$ is equal to the measure of EID within $\mathcal{N}$), and one representing the energy cost (Methods). (**B**) Same as **A**, but with the entropy minimization strategy. Unlike energy-constrained proportional betting, entropy minimization locally minimizes the entropy (maximizes the EID) of the posterior belief about the target location, in this case leading it to the nearby distractor by following the gradient. (**C–E**) An illustration of an animal tracking an object constrained to move in a line, in this case a hypothetical moth following a flower swaying in a breeze in a manner approximated by a 1-D sinusoid—a natural behavior reported in (***Sponberg et al., 2015***; ***Stockl et al., 2017***). We simulate the tracking of the flower using energy-constrained proportional betting. (**C**) In the top panel, we show an idealization of the moth's belief (blue line Gaussian distribution above the moth) about the flower's location when the flower reaches the center point. Higher values in the y direction represents higher confidence of the target at the given x location. The corresponding EID is overlaid in magenta; a darker color indicates higher expected information should the animal take a new sensory measurement in the corresponding location. Note the bimodal structure of the EID, with identical maximal information peaks on both sides of the Gaussian belief. (**D**) Simulation of the moth's position (red curve) while tracking a moving target (black curve) under a weak signal condition (Methods). The corresponding EID is overlaid in magenta. Note that even though trajectory segments are planned at 14 fixed-time intervals T (Methods) over the shown duration based on the EID at start of those intervals, the EID is here plotted continuously for visualization purposes only. Energy-constrained proportional betting results in persistent activation of movement when the EID is relatively diffuse due to lack of information. Note the presence of a distractor (marked by the black arrow) in the EID due to higher uncertainty in the sensory input as a result of the weaker signal. EIH responded to the distractor by making a detour away from the actual target position to gamble on the chance of acquiring more information, but does not get trapped by the distractor as in **B** because of the proportional betting strategy. (**E**) Same as (D), but under strong signal conditions. Now the energy-constrained proportional betting trajectory samples both peaks of the bimodal EID with excursions away from these peaks, similar to measured moth behavior (***Stockl et al., 2017***) (see Fig. 3G).

Results

Using one dataset we collected (electric fish), and three other previously published datasets, we performed side-by-side comparisons between one-dimensional trajectories generated by EIH and those measured from live animals during searching and tracking phases of movement in the pres-

ence of a target. Live animal data were either collected as one-dimensional trajectories, or collected as two-dimensional trajectories and projected to one dimension (Methods). Our data is from four species including South American gymnotid electric fish (glass knifefish *Eigenmannia virescens*, Valenciennes 1836) tracking a moving object in the dark under varying levels of electrical jamming; reanalyzed data from blind eastern American moles (*Scalopus aquaticus*, Linnaeus 1758) finding an odor source (***Catania, 2013***); reanalzyed data from American cockroach (*Periplaneta americana*, Linnaeus 1758) tracking an odor (***Lockey and Willis, 2015***); and reanalyzed data from hummingbird hawkmoth (*Macroglossum stellatarum*, Linnaeus 1758) tracking a robotically oscillated artificial flower while feeding (***Stockl et al., 2017***). Fig. S8 also shows an analysis of odor tracking in rats (***Khan et al., 2012***) but excluded from our analysis due to an insufficient number of trials. In these behavioral experiments, animals tracked or localized a signal source that was either stationary (mole and cockroach) or moving (electric fish and moth). Depending on the trial conditions, the live animal behavior trials were placed into one of two groups: a strong signal condition, and a weak signal condition (Methods).

Each of the live animal behavior datasets include experiments where the signal level of one sensory modality—considered to be important or dominant in driving tracking behavior—was varied. For all simulations with EIH, this one modality was selected for modeling. Each sensory system was modeled as a 1D point-sensor with a Gaussian model of observation that relates the sensory signal value to the variable that the animal is trying to estimate (*i.e.*, target location in tracking tasks, Methods). Sensor measurements were simulated by drawing values from the Gaussian model of observation. We added Gaussian measurement noise with variance determined by a specified signal-to-noise ratio (SNR) (Methods) to simulate the strong and weak signal conditions present in the live animal trials. These simulated measurements were used to update a probability distribution (often multi-peaked, such as Fig. 2D at arrow) representing the simulated animal's belief (Fig. 2C) about the target's likely location in the domain through a Bayesian update (***Thrun et al., 2005***) of the previous estimate (Methods, Supplementary Movie 2).

Animals and EIH generate more exploratory movement with weak signals.

We first examined the animals' tracking behavior under strong and weak signal conditions. Weakly electric fish engage in a behavior called "refuge tracking" where they try to maintain their position inside a close-fitting open-ended enclosure—such as a plastic tube—even as that enclosure is translated forward and backward along the lengthwise axis of the fish (***Rose and Canfield, 1993***). It is a natural behavior for holding position within protective cover being swayed by water flow, such as vegetation or root masses, during the fish's inactive (diurnal) periods in the South American rivers in which they live (***Rose and Canfield, 1993***). Prior work has shown that these fish will engage in larger amplitude body movements as sensory input is degraded (***Stamper et al., 2012***; ***Biswas et al., 2018***; ***Rose and Canfield, 1993***). For the trials reported here, all in the dark under infrared illumination, we degraded electrosensory input through varying the intensity of an externally imposed electrical jamming stimulus (Methods) which has previously been shown to impair electrolocation performance (***Watanabe and Takeda, 1963***; ***Bastian, 1987***; ***Ramcharitar et al., 2005***).

As shown in the first row of Fig. 3A, the weak signal resulted in more body movement. Throughout this work, we quantify the movement during tracking through a measure called "relative exploration," which is the amount of movement of the body divided by the minimum amount of movement required by perfect tracking (Methods). Across the full data set, we found a significant increase in relative exploration (Kruskal-Wallis test, $p < 0.001$, $n = 21$). The first row of Fig. 3B shows a representative trial of a mole engaging in a stationary odor source localization task (Methods). The behavioral data shows that the mole executes trajectories with significantly larger lateral oscillations under weak signal conditions (normal olfaction degraded by nostril blocking or crossing bilateral airflow (***Catania, 2013***)) as summarized in the relative exploration plot (Kruskal-Wallis test, $p < 0.001$, $n = 18$). Fig. 3C shows a typical strong and weak signal trial (degraded by trimming the olfactory antennae length, ***Lockey and Willis*** (***2015***)) of a cockroach localizing an odor source (Meth-

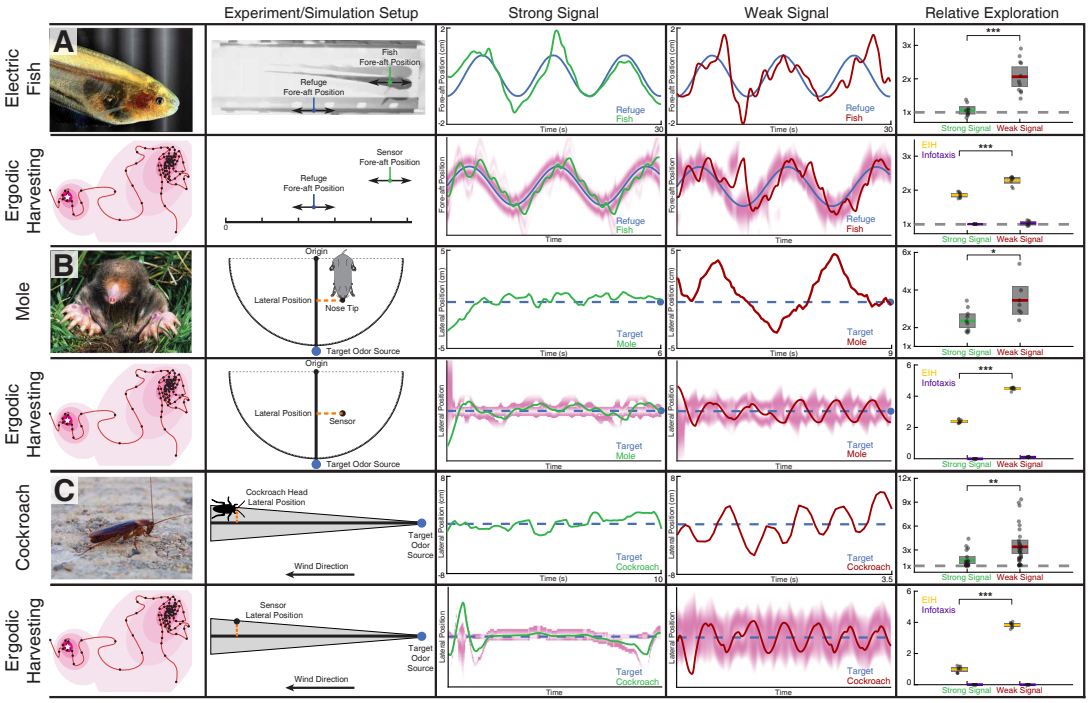

Figure 3. Representative trajectories of animals tracking a target compared to EIH, and relative exploration across all trials. Three representative live animal trajectories above trajectories generated by the EIH algorithm, length cropped for visual clarity. The moth data is not shown here due to the complexity of the prescribed target motion, but is shown in a subsequent figure. All EIH simulations were conducted with the same target path as present in the live animal data, using a signal level corresponding to the weak or strong signal categories (Methods). The EID is in magenta. Relative exploration across weak and strong signal trials (dots) shown in the right-most column (solid line: mean, fill is 95% confidence intervals). (**A**) The electric fish's fore-aft position was measured as it tracked a sinusoidally moving refuge. EIH (yellow) shows good agreement with behavioral data as both have significantly higher relative exploration when signal is weak (Kruskal-Wallis test, $p < 0.001$, $n = 21$ for experimental data, and $p < 0.001$, $n = 18$ for simulation). Infotaxis (purple), in contrast, leads to hugging the edge of the EID, resulting in smooth pursuit behavior as indicated by the near 1x relative exploration. (**B**) The experimental setup and data for the mole was extracted from a prior study (***Catania, 2013***). During the mole's approach to a stationary odor source, its lateral position with respect to the reference vector (from the origin to the target) was measured. Relative exploration was significantly higher under the weak signal conditions (Kruskal-Wallis test, $p < 0.012$, $n = 17$). In the second row, raw exploration data (lateral distance traveled in normalized simulation workspace units) are shown to allow comparison to simulation, as these are done in 1D (Methods). EIH shows good agreement with significantly increased exploration for weak signal (Kruskal-Wallis test, $p < 0.001$, $n = 18$), while infotaxis leads to cessation of movement. (**C**) The experimental setup and data for the cockroach was extracted from a prior study (***Lockey and Willis, 2015***). The cockroach head's lateral position was tracked and total travel distance was measured during the odor source localization task. Relative exploration is significantly higher for the weak signal condition (Kruskal-Wallis test, $p < 0.002$, $n = 51$). In the second row, we show that EIH raw exploration (as defined above) agrees well with measurements as the amount of exploration increased significantly under weak signal conditions (Kruskal-Wallis test, $p < 0.001$, $n = 18$), while infotaxis leads to cessation of movement. Asterisks indicate the range of p values for the Kruskal-Wallis test (* for $p < 0.05$, ** for $p < 0.01$, and *** for $p < 0.001$).

ods). Trials under weak signal conditions show an increased zig-zag amplitude, which leads to a significant increase in relative exploration in the weak signal condition as summarized in the plot (Kruskal-Wallis test, $p < 0.002$, $n = 51$). For the moth robotic-flower tracking behavior, the flower was moved with a sum-of-sine stimulus that cannot be visualized in the same manner as our first three species. We discuss the analysis of these trials in a following section.

In the second row of Fig. 3A-C, we show the corresponding EIH simulation when given the same target trajectory as used with the live animal trajectory, under simulated strong and weak signal

conditions. In these simulations, although the simulation has the target location provided, the ergodic information harvesting algorithm to track objects does not know this location and is only given simulated measurements (Algorithm S1). Overall, EIH shows good agreement with the measured animal behavior, with significantly increased relative exploration as the signal becomes weak (Kruskal-Wallis test, $p < 0.001$, $n = 18$ for each species).

Comparison to Infotaxis.

To assess the performance of infotaxis (***Vergassola et al., 2007***) across the species examined here, we used the same live animal target trajectories used for simulating EIH response as input. We computed relative exploration, and show this alongside EIH in the second row for each species in Fig. 3 (Methods). Infotaxis leads to mostly smooth tracking trajectories with near unity relative exploration when target is moving (Fig. 3A); there is no increase in exploration with reduction of signal strength as in the original behavior and EIH simulations. When the target is stationary, infotaxis leads to cessation of movement, as shown in Fig. 3B-C. This change in the tracking response is further analyzed in Fig. S3.

Increased exploratory movement and sensing-related high frequency movement.

To further characterize the sensing-related movement patterns and verify whether the increase in exploration is mainly due to these movements, we performed a spectral analysis of the animal's tracking response. In Fig. 4A-B we show the frequency spectrum of weakly electric fish refuge tracking behavior under strong and weak signal conditions. Two frequency bands can be identified: 1) a baseline tracking band that overlaps with the frequency at which the target (the refuge) was moved; and 2) a sensing-related movement frequency band that accounts for most of the increased exploratory movements as signal weakens. The Fourier magnitude is significantly higher for the sensing-related movement frequency band under weak signal conditions when compared with strong signal conditions, for both the measured animal behavior (Fig. 4C, Kruskal-Wallis test, $p < 0.001$) and EIH simulations (Fig. 4D, Kruskal-Wallis test, $p < 0.001$). For animals localizing a stationary target (Fig. 4E-L), the whole spectrum is considered to be sensing-related, since the target is stationary. Similar to electric fish, we found significantly increased mean Fourier magnitude within the sensing-related frequency band for mole (Fig. 4G-H, Kruskal-Wallis test, $p < 0.009$ for measured mole response and $p < 0.001$ for simulation) and cockroach (Fig. 4K-L, Kruskal-Wallis test, $p < 0.003$ for measured cockroach response and $p < 0.001$ for simulation) when compared to strong signal conditions. These results confirm that the increased exploratory movement shown earlier is related to sensory organ wiggles, as predicted by EIH.

For evaluation of the moth data where the target moves in a more complex sum-of-sines stimulus, we performed a similar spectral analysis when the moth was tracking the robotic flower under high illumination (strong signal) and low illumination (weak signal) conditions Fig. 4M-P (***Stockl et al., 2017***). We analyzed the first 18 prime frequency components (up to 13.7 Hz) of both the moth's response (Fig. 4M, data from (***Stockl et al., 2017***), $n = 13$ for strong signal and $n = 10$ for weak signal) and simulation (Fig. 4N, $n = 120$ for strong signal and $n = 120$ for weak signal), which is the same range used in (***Stockl et al., 2017***). We show the spectrum in Fig. 4M-N as a Bode gain plot rather than Fourier magnitude since the target spectrum now covers a wide frequency band including sensing-related movements (Methods). Consistent with previously reported behavior (***Stockl et al., 2017***), we found significantly increased mean tracking gain in the moth's response within the mid-range frequency region relative to the strong signal condition (Fig. 4O, Kruskal-Wallis test, $p < 0.02$, $n = 23$). This pattern is predicted by EIH simulations with the same sum-of-sine target trajectory (Fig. 4P, Kruskal-Wallis test, $p < 0.001$, $n = 240$).

Sensing-related movements increase fish refuge tracking performance.

A crucial issue to address is whether the additional sensing-related motions found in weak signal conditions in animals, and predicted by EIH, cause increased tracking performance. To answer this question, we constructed a filter to selectively attenuate only the higher frequency motion

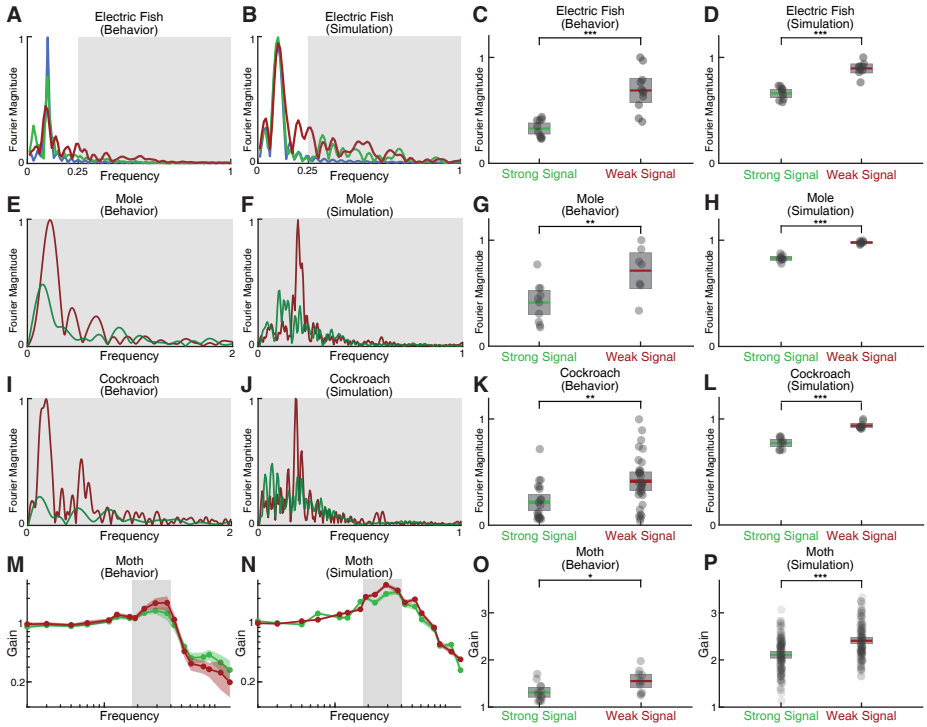

Figure 4. Spectral analysis of live animal behavior and simulated behavior. All the single trial Fourier spectra shown in A-B, E-F, and I-J are for the trials shown in Fig. 3. (**A-B**) Representative Fourier spectra of the measured and modeled (EIH-generated) refuge tracking trajectories, with target trajectory (blue), strong signal condition (green), and weak signal condition (red). The frequencies above the frequency domain of the target's movement is shaded; components in the non-shaded region are excluded in the subsequent analysis. The Fourier magnitude is normalized. (**C-D**) Distribution of the mean Fourier magnitude within the sensing-related movement frequency band (gray shaded region) for strong and weak signal trials. Each dot represents a behavior trial or simulation. Significantly higher magnitude is found within the sensing-related movement band under weak signal conditions (Kruskal–Wallis test, $p < 0.001$, $n = 21$ for measured behavior, and $p < 0.001$, $n = 18$ for EIH). (**E-F**) Representative Fourier spectra of the measured and modeled head movement of a mole while searching for an odor source (lateral motion only, transverse to the line between origin and target), plotted as in **A-B**. Because the target is stationary in contrast to the fish data, the entire frequency spectrum was analyzed. (**G-H**) Mean Fourier magnitude distribution. In the weak signal condition, there is significantly more lateral movement power under weak signal in both the animal data and simulation (Kruskal–Wallis test, $p < 0.009$, $n = 17$ for behavior data, and $p < 0.001$, $n = 18$ for EIH simulations). (**I-J**) Fourier spectra of mole's lateral trajectory, plotted the same as in **A-B**. (**K-L**) Mean Fourier magnitude distribution. Significantly higher power is found under weak signal in the behavior data and simulations (Kruskal–Wallis test, $p < 0.003$, $n = 51$ for behavior data, and $p < 0.001$, $n = 18$ for EIH simulations). (**M-N**) Bode magnitude plot (Methods) of a moth tracking a robotic flower that moves in a sum-of-sine trajectory (figure reprinted from ***Stockl et al.*** (***2017***)) and the corresponding simulation using the same target trajectory. Each dot in the Bode plot indicates a decomposed frequency sample from the first 18 prime harmonic frequency components of the flower's sum-of-sine trajectory. A total of 23 trials ($n = 13$ for strong signal, and $n = 10$ for weak signal) were used to establish the 95% confidence interval shown in the colored region. Note the increase in gain in the midrange frequency region (shaded in gray) between the strong signal and weak signal conditions. The confidence interval for the simulation Bode plot is established through 240 trials ($n = 120$ for strong signal, and $n = 120$ for weak signal). The same midrange frequency regions (shaded in gray) are used for the subsequent analysis. (**O-P**) Mean midrange frequency tracking gain distribution for the moth behavior and simulation trials. Moth's exhibit a significant increase in midrange tracking gain for the weaker signal condition (Kruskal–Wallis test, $p < 0.02$, $n = 23$), in good agreement with simulation (Kruskal–Wallis test, $p < 0.001$, $n = 240$). Asterisks indicate the range of p values for the Kruskal-Wallis test.

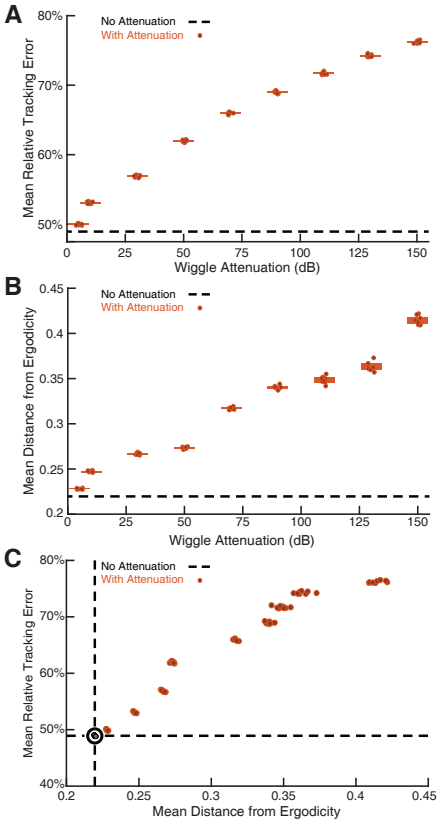

Figure 5. The increase in full-body oscillations during refuge tracking in electric fish reduces tracking error. The full-body oscillation in the simulated sensor trajectory from Fig. 3A (EIH under weak signal) was gradually removed at increasing wiggle attenuations (Methods) with 8 trials per condition (total $n = 80$) to establish the confidence interval. (**A**) Relative tracking error (% of amplitude of fore-aft movement, Methods) as sensor oscillation is attenuated. The line near 45% error shows relative tracking error for the original unfiltered trajectory (0 dB attenuation). The thickness of the horizontal bars represents the 95% confidence interval across the eight trials (individual dots) for each attenuation condition. For each attenuation levels, the individual trial dots are plotted with small horizontal offset to enhance clarity. The baseline tracking error with intact wiggle is marked by the dashed black line. (**B**) Distance from ergodicity (Methods) as a function of wiggle attenuation. Zero distance from ergodicity indicates the optimal trajectory that perfectly matches the statistics of the EID. As the distance increases from zero, the corresponding trajectory becomes less and and less related to the EID. The baseline data with intact wiggle is marked by the dashed black line. (**C**) Relative tracking error plotted against distance from ergodicity across wiggle attenuations. There is a clear positive correlation between tracking error and distance from ergodicity as wiggle is progressively attenuated.

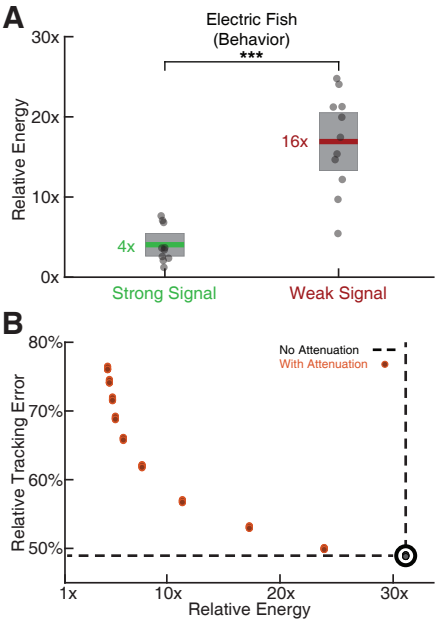

Figure 6. Relative energy for electric fish tracking behavior and EIH simulation with attenuated body oscillations. (**A**) Relative energy used by the electric fish during refuge tracking behavior under strong and weak signal conditions. Trials are the same as shown in Fig. 3A. Relative energy is the ratio of the mechanical energy needed to move the fish along the tracking trajectory divided by the energy to move the fish along the refuge trajectory. Weak signal conditions show a significantly higher relative energy as a result of the additional sensing-related movements (Kruskal-Wallis test, $p < 0.001$, $n = 21$). (**B**) Relative energy and relative tracking error for EIH simulations as sensing-related movements are progressively attenuated (data from Fig. 5A-C, weak signal condition). As the wiggles are attenuated, simulating investment of less mechanical effort for tracking, the relative energy decreases from 30x to less than 5x, but tracking error increases from 50% to around 75%. This plot also suggests a possible diminishing return in tracking error with additional energy expenditure past 30x. The lower bound near 4x is similar to the relative energy for the strong signal condition, and arises due to small disparaties from the perfect sinusoid that the refuge is following (the 1x path). Asterisks indicate the range of p values for the Kruskal-Wallis test.

components without affecting the baseline tracking motion (Methods). Simulated weakly electric fish tracking trajectories in the weak signal condition—similar to that shown in the second row of Fig. 3A "Weak Signal"—were filtered at increasingly high levels of attenuation (Fig. 4B, shaded region). This led to a visible decrease in sensing-related body oscillations (Fig. S7). These filtered trajectories were then provided as the input to a sinusoidal tracking simulation in which the sensor moved according to the filtered trajectory to take measurements and update an initially uniform belief in the same way as EIH (Methods).

We show the results in Fig. 5A in terms of relative tracking error, where 50% error means a departure from perfect tracking that is one-half the amplitude of the fore-aft sinusoidal motion. Relative tracking error increases in proportion to the amount of sensing-related motion attenuation, from ≈50% with no attenuation to ≈75% with the highest attenuation we used. We then evaluated the distance from ergodicity, a dimensionless quantity that measures how well a given trajectory matches the optimal proportional betting trajectory (Methods), for all the trajectories. We found that an increase in attenuation also leads to monotonically increased distance from ergodicity. This indicates that the filtered trajectories are progressively worse at proportionally betting on information (Fig. 5B). Fig. 5C combines these two analyses, demonstrating that the distance from ergodicity is proportionate to tracking error.

Error versus energy expenditure during fish refuge tracking. We estimated the mechanical energy needed to move the fish body along the measured trajectories in comparison to moving the body along the exact trajectory of the refuge. The result is computed as relative energy by

taking the ratio of energy needed for moving along the measured tracking trajectory to the energy needed for moving the body along the refuge's trajectory (Methods). Electric fish spent significantly more energy during tracking in the weak signal condition compared to the strong signal condition (≈ 4x more, Fig. 6A, Kruskal-Wallis test, $p < 0.001$, $n = 21$). Finally, we examined how tracking error relates to the energy expended on moving the body, starting with unfiltered (EIH) trajectories and then through the series of stepped attenuation trajectories that progressively eliminate sensing-related body oscillations. This was done by computing the relative energy for the simulation data shown in Fig. 5A. We found that the tracking error decreased as the relative energy increased, with diminishing returns as the relative energy level nears that needed for the original unfiltered EIH trajectory (≈ 30 times the energy needed to move the body along the refuge trajectory, Fig. 6B).

Discussion

The body's mechanical and information processing systems have coevolved to afford behaviors that enhance evolutionary fitness. Shortly after Shannon published his work on the information capacity of communication channels (***Shannon and Weaver, 1949***), his ideas were applied to visual perception (***Attneave, 1954***; ***Barlow, 1959***) to describe ideas of efficient coding in the visual periphery. Since then, continual progress has been made in applying information theory to illuminate a host of problems on the coding and energetics of sensory signals from receptors through to central nervous system processing (***Atick, 1992***; ***Laughlin et al., 1998***; ***Niven and Laughlin, 2008***; ***Sengupta et al., 2010***). A parallel literature has matured analyzing animal motion (***Waldron et al., 2008***; ***Srinivasan and Ruina, 2006***; ***Ramdya et al., 2017***; ***Nyakatura et al., 2019***; ***Aguilar et al., 2016***; ***Collins et al., 2005***; ***Lee et al., 2008***; ***McInroe et al., 2016***; ***Sefati et al., 2013***). A domain of increasing interest is connecting the information gathered through movement to the analysis of movement (***Cowan and Fortune, 2007***; ***Rucci and Victor, 2015***; ***MacIver et al., 2010***; ***Sprayberry and Daniel, 2007***; ***Stamper et al., 2012***; ***Biswas et al., 2018***; ***Yovel et al., 2010***; ***Stockl et al., 2017***; ***Fujioka et al., 2016***; ***Yovel et al., 2011***; ***Ghose and Moss, 2003***, ***2006***; ***Hofmann et al., 2014***; ***Bar et al., 2015***; ***Nelson and MacIver, 2006***; ***Bush et al., 2016***; ***Yang et al., 2016***), but a generalized framework to bridge the gap between information gained over sense organ or whole body trajectories and the energetics of following such trajectories is lacking. EIH is a candidate framework that is sufficiently general to invite application to a host of information-related movements observed in living organisms, while sufficiently well-specified to generate testable quantitative predictions.

In the insects-to-mammals assemblage of animal species analyzed above, we observe gambling on information through motion, where the magnitude of the gamble is indexed by the energy it requires. EIH's approach of extremizing ergodicity generates trajectories that bet on information, exchanging units of energy for the opportunity to obtain a measurement in a new high-value location. For both the measured tracking trials and their EIH simulated versions, a key change that occurs as sensory signals weaken is an increase in the rate and amplitude of the excursions from the mean trajectory, which we have quantified as an increase in relative exploration. We do not presently have a complete mechanistic understanding of the cause of these motions in animals, but we can interrogate how they arise within EIH. First, the increase in the size of these excursions in weak signal conditions arises because the EID spreads out in these conditions due to high uncertainty (for example, wider magenta bands in Fig. 3). Since EIH samples proportionate to the expected information, as information diffuses, the excursions to sample the more spread-out EID will increase in size.

Second, in EIH the frequency of these excursions is related to how far ahead in time a trajectory is generated for (variable T in Algorithm S1). One can consider this analogous to how far ahead an animal generates a trajectory for before a change in information can result in a change in their trajectory. For example, tiger beetles see their prey, then execute a trajectory to the prey that is completed regardless of subsequent motion of the prey (***Gilbert, 1997***). In EIH, over the course of the generated trajectory epoch, changes in the expected information density due to new sensory observations similarly have no effect; these will only be incorporated for the next trajectory

segment.

In EIH, perfect ergodicity is approached through the trade-off between the ergodic metric and control effort within the prescribed generated trajectory time horizon T. As T asymptotically approaches infinite duration, the system will approach perfect ergodicity as the ergodic measure approaches zero. Conversely, as T asymptotically approaches zero time duration, the environment will be minimally sampled and EIH executes the best trade-off between the cost of movement and extremizing ergodicity in that minimal segment of time. Changing T between these bounds will affect the frequency of the sensing-related excursions. This is because the information landscape is assumed to be static until T has elapsed, the EID is updated, and a new segment of trajectory is generated. For example, imagine a typical simulated one-dimensional trajectory that exhibits wiggle motions. As the simulated animal moves to visit a region of dense expected information in one direction, the unvisited locations in the other direction start to accumulate uncertainty in the belief (at a rate proportional to the noise level), leading to an increase in the density of expected information in those locations (Fig. S2). After the sensor finishes the current trajectory segment of duration T, it then moves in the opposite direction to explore the unvisited regions with a high expected information density (Fig. S2). Thus, a shorter T causes the sensor to react more quickly in response to changes in the information landscape and hence to a higher wiggle frequency. The initial T (see Table S1) used for the behavior simulations was chosen to fit the frequency of sensing-related oscillations observed in the weakly electric fish refuge tracking data. The same value was applied to mole and cockroach trials, and reduced by a factor of five for the moth data due to the higher frequency content of the prescribed robotic flower movement.

As gambling on information through motion involves a trade-off between increasing how well a trajectory approaches ideal sampling (ergodicity) and reducing energy expenditure, a useful quantity to examine is how tracking error changes with the energy expended on motion. To do so, we estimated the mechanical energy needed to move the body of the electric fish along the weak and strong signal trajectories, and found that weak signal trajectories required four times as much energy to move the body along as strong signal trajectories. In simulation, we examined how tracking error changes as more energy is invested in the sensing-related movements. This analysis shows that the accuracy of tracking increases with the mechanical effort expended on the small whole-body oscillations that these fish make while tracking, with a 25% reduction of tracking error at the highest level of energy expenditure compared to the low energy case where sensing related movements are removed.

Comparison to information maximization. Information maximization and EIH emphasize different factors in target tracking. First, if a scene is so noisy as to have illusory targets (more than one peak in the probability distribution representing the estimated target location, called the belief), or *actually* includes multiple objects, information maximization will not result in sensing the information distributed across the scene, but rather move until a local information maxima is reached (for example, the distractor in Fig. 2B) and stay at that location. With energy-constrained proportional betting, information across a specified region of interest will be sampled in proportion to its expected magnitude (Fig. 2A). This leads to sensing-related movements that may, at first glance, seem poorly suited to the task: for example, if the distractor has higher information density, as it does in Fig. 2A, then it will be sampled more often than the lower information density of the true target—but what is important here is that the true target is sampled at all, enabling the animal to avoid getting trapped in the local information maxima of the distractor. For information maximization, if 1) there is only one target of interest; 2) the EID is normally distributed; and 3) the signal is strong enough that false positives or other unmodeled uncertainties will not arise, then information maximization will reduce the variance of the estimated location of the single target being sought and direct movement toward the true target location. We interpret the poor agreement of simulated information maximization trajectories with measured behavior in our results as indicating that the conjunction of these conditions rarely occur in the animal behaviors we examine.

The second area where these two approaches have different emphases is highlighted in cases

where noise is dominating sensory input in high uncertainty scenarios as is common in naturalistic cases. Information maximization leads to a cessation of movement since no additional information is gained in moving from the current location (Fig. S1C, Fig. S3A). Energy-constrained proportional betting will result in a trajectory which covers the space (Fig. S1C, Fig. S3A): the expected information is flat, and a trajectory matching those statistics is one sweeping over the majority of the workspace at a density constrained by EIH's balancing of ergodicity with energy expenditure. For information maximization, coverage can only be an accidental byproduct of motions driven by information maximization.

Other interpretations of the behavioral findings. Yovel et al. (***Yovel et al., 2010***) suggested that off-axis sensing behavior arises from moving to a peak in the EID (similar to Fig. 2B). This hypothesis is similar to infotaxis in terms of using information maximization, although using Fisher Information rather than entropy minimization. Yovel et al. show in simulation that the information maximization strategy leads to a smooth tracking trajectory which hugs the edge of the signal trail (at one of the two information peaks, Fig. 4 of (***Yovel et al., 2010***)). This is replicated by our simulation of infotaxis as well (Fig. S1D). Figure 2C of their paper shows the average intensity of the emitted sound as a function of azimuth to target, but does not show the distribution at each azimuth. As a result the slope of this curve appears to be uniquely determined and centered on the target, whereas statistical analysis would indicate a range of maximal slope locations that is dependent on the distribution of intensity at each azimuth. Hence, their analysis does not distinguish between information maximization and proportional betting on information because the slopes of emitted sound intensity could mean that the actual optimizer is over a potentially large range of possibilities. Further, hugging one edge of the EID deviates from our results and another study examining rat odor tracking (***Khan et al., 2012***). ***Khan et al.*** (***2012***) show that in rat odor tracking behavior only about 12% of the trajectory qualifies as edge-tracking, suggesting that the rat's zig-zagging trajectory is not centered on the edge of the trail—as predicted by the information maximization hypothesis—but rather on the middle of the odor trail, consistent with ergodic harvesting.

Another hypothesis is that active sensing movements arise from the animal adapting its closed-loop tracking gain response to a reduction in signal contrast (***Borst et al., 2005***; ***Ghose and Moss, 2006***; ***Maimon et al., 2010***; ***Biswas et al., 2018***). However, this gain adaptation hypothesis is underspecified, in the sense that critical components are missing to formulate an algorithm that generates predictive trajectories conforming to the hypothesis. If gain adaptation is implemented with a Bayesian filter and a process is specified to generate oscillatory motion around targets according to the variance of the belief as a measure of uncertainty, then in the narrow context of a single target with no distractors (neither real nor fictive due to high uncertainty), such an algorithm can be tuned to behave similarly to EIH. However, in more realistic scenarios, there is no apparent mechanism to address real or fictive distractors, a capability of EIH we elaborate on further below. Further work is needed to test the differences between EIH and the gain adaptation hypothesis, or to determine whether gain adaptation is an implementation of EIH in specific, biologically relevant circumstances.

Khan et al. (***2012***) introduced a model for odor tracking that instructs the sensor to move forward and lateral at a fixed velocity and make decisions to switch the direction of lateral movements based on specific events of sensor measurement and position. Though their model could in principle be tweaked to fit the trajectories of animal tracking under weak signal, the zig-zag wiggle movements are explicitly programmed to appear based on ad-hoc strategies. This makes the model less generalizable and yields little insight on the underlying mechanism. In contrast, the wiggle movements that emerge with EIH are not programmed but arise naturally from the need to provide coverage on information. This view is supported by the wiggle attenuation analysis shown in Fig. 5. In addition, the Khan et al. model lacks the ability to address distractors, as shown in Fig. S2 and Fig. S5, since the movement strategy is not based on the belief or EID map, whereas EIH naturally provides coverage in these scenarios.

Finally, Rucci and Victor (***Rucci and Victor, 2015***) and Stamper et al. (***Stamper et al., 2012***) propose that active sensing movements are the outcome of an animal actively matching the spatial-temporal dynamics of upstream neural processing—a process by which the movement serves as a "whitening filter" (***Rucci and Victor, 2015***) or "high-pass filter" (***Stamper et al., 2012***). Sensing-related movements could be for preventing perceptual fading (***Kunapareddy and Cowan, 2018***), which has similarities to the high pass filter hypothesis in that motion is to counter sensory adaptation, a high pass filter-like phenomena. Although evidence for the perceptual fading hypothesis during tracking behaviors is lacking, EIH shows good agreement with animal behavior without any mechanism for sensory adaptation. Similar to the gain adaptation hypothesis, the high-pass filter hypothesis is also missing key components for trajectory prediction. Nonetheless, when implemented with the missing components, including a Bayesian filter and a feedback process that generates trajectories that match the desired spatial-temporal dynamics (***Biswas et al., 2018***), the high-pass filter hypothesis does not conflict with EIH in single target cases with low uncertainty. This is because EIH also predicts a preferred frequency band for wiggle movements that may match the preferred spectral power of upstream neural processing. However, in the context of multiple target scenarios, high uncertainty due to weak signal resulting in fictive distractors, or in the absence of any target, the same considerations apply to the high-pass filter hypothesis as were mentioned for the gain adaptation hypothesis.

Distractors and multiple targets. Given the above discussion, a capability of EIH that differentiates it from prior theories and that naturally arises from its distributed sampling approach is its ability to reject distractors and sample multiple targets. The live animal experimental data we analyzed did not feature either real distractors (here defined as objects having a distinguishably different observation model from that of the target) or multiple targets (multiple objects with identical observation models). Nonetheless, the EIH simulations disclose that what we are calling "sensing-related motions"—those movements that increase as sensory signals weaken—sometimes occur for rejection of fictive distractors. A fictive distractor emerges when the current belief for the target's location becomes multi-peaked; each peak away from the true target's location is then a fictive distractor (illustrated by arrow in Fig. 2D). Fig. S2 shows the presence of these fictive distractors in the simulations of the fish, cockroach, and mole tracking behaviors, where we plot the belief rather than the EID. Fictive distractors arise in both the strong and weak signal conditions, but result in small amplitude excursions in the strong signal conditions because of the higher confidence of observations. We hypothesize that a similar process of fictive distractor rejection is one source of the increase in sensing-related wiggle movements in the animal data as sensory signals weaken. In the simulated tracking behavior, the other source of sensing-related movements is the increased spread of a (unimodal) EID as signals weaken, as earlier discussed.

False positive rejection has the signature of a digression from the nominal tracking trajectory; this digression ends when one or more samples have been received indicating there is no object present at the spurious belief peak, which then brings the believed target location back to somewhere closer to the true target position (Fig. S2). In contrast, with a physical distractor, a digression should occur, but the observations support that the object being detected has a different observation model from that of the target, rather than the absence of an object. As none of our datasets include physical distractors, we investigated EIH's behavior in this case with a simulated physical distractor. Fig. S5 shows a simulated stationary physical distractor in addition to a stationary target. EIH is able to locate the desired target while rejecting the distractor. This result buttresses a finding in a prior robotics study, where we experimentally tested how EIH responded to the presence of a physical distractor and showed that an electrolocation robot initially sampled the distractor but eventually rejected it (Fig. 8 of ***Miller et al. 2016***). In comparison, Fig. S5 shows that infotaxis stalls as it gets trapped at one of the information maximizing peaks and fails to reject the distractor.

If, instead of a distractor and a target, EIH has two targets, the advantage of EIH's sampling the workspace proportional to the information density is particularly well highlighted. A simulation of this condition is shown in Fig. S4. EIH maintains good tracking with an oscillatory motion providing

coverage for both of the targets. As seen in Fig. S4, such coverage is not a feature of infotaxis, which gets stuck at the location of the first target and fails to detect the presence of the other target. A final case to consider is multiple targets with different (rather than identical) observation models. Tracking in such cases requires a simple adjustment to the calculation of the EID that we have explored elsewhere (Eq. 13 of ***Miller et al. 2016***).

While these preliminary simulations exploring how EIH performs with multiple targets and distractors are promising, it points to a clear need for animal tracking data in the presence of physical distractors or multiple targets (and in 2-D or 3-D behaviors: Appendix 6) in order to better understand whether EIH predicts sensory organ motion better than the gain adapatation or high-pass filtering theories in these cases.

Biological implementation. The sensor or whole-body wiggle we observe in our results is for proportional betting with regard to sensory system-specific EIDs—for electrosense, olfaction, and vision. To implement EIH, one needs to store at least a belief encoding knowledge about the target. The Bayesian filter update in EIH has the Markovian property, meaning that only the most recent belief is required to be stored. The EID, moreover, is derived from the belief and only used for every generated trajectory segment update, hence does not need to be stored. While the memory needs of EIH are low, trajectory synthesis requires computing the distance from ergodicity between candidate trajectories and the EID, a potentially complex calculation (Methods, Appendix 3). However, the complexity of our calculation may not be indicative of the complexity of implementation in biology. For instance, a recent study (***Stachenfeld et al., 2017***) suggests that a predictive map of future state is encoded in grid cells of the entorhinal cortex through spatial decomposition on a low-dimensionality basis set—a process similar to the calculation of the ergodic metric (Appendix 3). In weakly electric fish, electroreceptor afferents have power law adaptation in their firing rate in response to sensory input (***Drew and Abbott, 2006***; ***Clarke et al., 2013***). This makes their response invariant to the speed of the target (***Clarke et al., 2013***) and hence similar to the simulated sensory input used to drive EIH. The power law adaptation also results in a very strong response at the reversal point during whole-body oscillations (Fig. 5 of ***Clarke et al. 2014***), a response generated by hindbrain-midbrain feedback loops (***Clarke and Maler, 2017***). Given the importance of an increased rate of reversals as signals become weaker in the fish tracking data and EIH, and EIH's invariance to speed, the hindbrain area along with feedback loops to the midbrain are a worthwhile target for future work on the biological basis of EIH.

Ergodic movement as an embodied component of information processing. As in the case of fixational eye movements (***Rucci and Victor, 2015***), a common interpretation of body or sensor organ movements away from the assumed singular goal trajectory is that this reflects noise in perceptual or motor processes. Ergodic harvesting presents a competing hypothesis: gathering information in complex environments means the system should be observed to move away from the singular goal trajectory. These excursions occur as predicted by EIH, including the possibility of multiple targets, and thus increase in size when that information is more diffuse.

If an animal is at one peak of a multi-peaked belief distribution, what motivates it to move away from the current peak? The current peak already has sensor noise and other aspects of sensor physics incorporated, but misses other important sources of uncertainty. An occlusion may corrupt signal quality at a location otherwise predicted to have high target information. Other signal generators in the environment may emit confusing signals or the location may be contaminated by a fictive distractor arising from unmodeled uncertainty. Hence, the opportunity to visit another location in space that is statistically independent, yet contains a similar amount of predicted information, gives an animal an opportunity to mitigate unmodeled uncertainties through energetic expense. This is supported by experiments in human visual search suggesting that saccades are planned in a multi-stage manner for coverage of information towards the task-relevant goal rather than aiming for information maximization (***Yang et al., 2016***; ***Hoppe and Rothkopf, 2019***). For example, a model to predict human visual scan paths found 70% of the measured fixation locations were efficient from an information maximization perspective, but there were many fixations (≈30%) that

were not purely for maximizing information and attributed in part to perceptual or motor noise (***Yang et al., 2016***). We hypothesize that these apparently less efficient fixation locations are in fact the result of gambling on information through motion. It is also possible that motor noise may aid coverage in a computationally inexpensive manner.

The role of motion in this sensing setting is to mitigate the adverse impact of sensor properties. If, however, one is in an uncertain world full of surprises that cannot be anticipated, using energy to more fully measure the world's properties makes sense. This is like hunting for a particular target in a world where the environment has suddenly turned into a funhouse hall of mirrors. Just as finding one's way through a hall of mirrors involves many uses of the body as an information probe—ducking and weaving, and reaching out to touch surfaces—EIH predicts amplified energy expenditure in response to large structural uncertainties.

Acknowledgments

We thank Mark Willis, Simon Sponberg, and Ken Catania for providing the original behavioral tracking data used for the studies we have cited. We thank the anonymous reviewers for many improvements as well as a suggestion on biological implementation. We thank Madhav Mani and Brennan Sprinkle for helpful discussions and feedback on an earlier draft. Funded by National Science Foundation IIS-1427419.

Author Contributions

M.A.M., T.D.M., and C.C. designed research; C.C. performed research; T.D.M., C.C., and M.A.M. contributed new reagents/analytic tools; C.C. analyzed data; and M.A.M., T.D.M., and C.C. wrote the paper.

Competing Interest

The authors declare no competing financial interest.

References

Aguilar J, Zhang T, Qian F, Kingsbury M, McInroe B, Mazouchova N, Li C, Maladen R, Gong C, Travers M, Hatton RL, Choset H, Umbanhowar PB, Goldman DI. A review on locomotion robophysics: the study of movement at the intersection of robotics, soft matter and dynamical systems. Rep Prog Phys. 2016 sep;

Atick JJ. Could information theory provide an ecological theory of sensory processing? Network: Computation in Neural Systems. 1992;

Attneave F. Informational aspects of visual perception. Psychol Rev. 1954;

Bar NS, Skogestad S, Marçal JM, Ulanovsky N, Yovel Y. A Sensory-Motor Control Model of Animal Flight Explains Why Bats Fly Differently in Light Versus Dark. PLoS Biol. 2015 01;

Barlow HB. Sensory mechanisms, the reduction of redundancy, and intelligence. In: *NPL Symposium on the Mechanization of Thought Process* No. 10, HM Stationery Office, London; 1959.

Basil JA, Hanlon RT, Sheikh SI, Atema J. Three-dimensional odor tracking by Nautilus pompilius. J Exp Biol. 2000;

Bastian J. Electrolocation in the presence of jamming signals: behavior. J Comp Physiol A. 1987;

Biswas D, Arend LA, Stamper SA, Vágvölgyi BP, Fortune ES, Cowan NJ. Closed-Loop Control of Active Sensing Movements Regulates Sensory Slip. Curr Biol. 2018;

Borst A, Flanagin VL, Sompolinsky H. Adaptation without parameter change: dynamic gain control in motion detection. Proceedings of the National Academy of Sciences. 2005;

Bush NE, Solla SA, Hartmann MJ. Whisking mechanics and active sensing. Curr Opin Neurol. 2016;

Calhoun AJ, Chalasani SH, Sharpee TO. Maximally informative foraging by Caenorhabditis elegans. Elife. 2014;

Caputi A, Aguilera P, Castelló M. Probability and amplitude of novelty responses as a function of the change in contrast of the reafferent image in G. carapo. J Exp Biol. 2003;

Catania KC. Stereo and serial sniffing guide navigation to an odour source in a mammal. Nat Commun. 2013;

Censi A, Straw AD, Sayaman RW, Murray RM, Dickinson MH. Discriminating External and Internal Causes for Heading Changes in Freely Flying Drosophila. PLOS Computational Biology. 2013 02;

Chen C, Murphey T, MacIver MA, Code and data to reproduce results from "Sense organ control in moths to moles is a gamble on information through motion". Github

Clarke SE, Longtin A, Maler L. The neural dynamics of sensory focus. Nat Commun. 2015;

Clarke SE, Longtin A, Maler L. A Neural Code for Looming and Receding Motion Is Distributed over a Population of Electrosensory ON and OFF Contrast Cells. Journal of Neuroscience. 2014;

Clarke SE, Maler L. Feedback Synthesizes Neural Codes for Motion. Current Biology. 2017;

Clarke SE, Naud R, Longtin A, Maler L. Speed-invariant encoding of looming object distance requires power law spike rate adaptation. Proceedings of the National Academy of Sciences. 2013;

Collins S, Ruina A, Tedrake R, Wisse M. Efficient bipedal robots based on passive-dynamic walkers. Science. 2005;

Cowan NJ, Fortune ES. The critical role of locomotion dynamics in decoding sensory systems. J Neurosci. 2007;

De La Torre G, Flaßkamp K, Prabhakar A, Murphey TD. Ergodic exploration with stochastic sensor dynamics. In: *2016 American Control Conference (ACC)* IEEE; 2016.

Drew PJ, Abbott LF. Models and Properties of Power-Law Adaptation in Neural Systems. Journal of Neurophysiology. 2006; PMID: 16641386.

Duistermars BJ, Chow DM, Frye MA. Flies require bilateral sensory input to track odor gradients in flight. Curr Biol. 2009 Aug;

Ferner MC, Weissburg MJ. Slow-moving predatory gastropods track prey odors in fast and turbulent flow. J Exp Biol. 2005;

Ferris BD, Green J, Maimon G. Abolishment of spontaneous flight turns in visually responsive Drosophila. Current Biology. 2018;

Fujioka E, Aihara I, Sumiya M, Aihara K, Hiryu S. Echolocating bats use future-target information for optimal foraging. Proc Natl Acad Sci USA. 2016 Apr;

Ghose K, Moss CF. The sonar beam pattern of a flying bat as it tracks tethered insects. J Acoust Soc Am. 2003;

Ghose K, Moss CF. Steering by Hearing: A Bat's Acoustic Gaze Is Linked to Its Flight Motor Output by a Delayed, Adaptive Linear Law. J Neurosci. 2006;

Gilbert C. Visual control of cursorial prey pursuit by tiger beetles (Cicindelidae). J Comp Physiol A. 1997;

Gottlieb J, Hayhoe M, Hikosaka O, Rangel A. Attention, reward, and information seeking. J Neurosci. 2014 Nov;

Hofmann V, Geurten BRH, Sanguinetti-Scheck JI, Gómez-Sena L, Engelmann J. Motor patterns during active electrosensory acquisition. Frontiers in Behavioral Neuroscience. 2014;

Hoppe D, Rothkopf CA. Multi-step planning of eye movements in visual search. Sci Rep. 2019;

Khan AG, Sarangi M, Bhalla US. Rats track odour trails accurately using a multi-layered strategy with near-optimal sampling. Nat Commun. 2012;

Kokoska S, Zwillinger D. CRC standard probability and statistics tables and formulae. Crc Press

Kording KP, Wolpert DM. Bayesian integration in sensorimotor learning. Nature. 2004 Jan;

Kothari NB, Wohlgemuth MJ, Moss CF. Adaptive sonar call timing supports target tracking in echolocating bats. J Exp Biol.

Kunapareddy A, Cowan NJ. Recovering observability via active sensing. In: *2018 Annual American Control Conference (ACC)* IEEE; 2018.

Lasota A, Mackey MC. Chaos, fractals, and noise: stochastic aspects of dynamics, vol. 97. Springer

Laughlin S, van Steveninck R, Anderson J. The metabolic cost of neural information. Nat Neurosci. 1998 MAY;

Lee J, Sponberg S, Y Loh O, Lamperski A, J Full R, Cowan N. Templates and Anchors for Antenna-Based Wall Following in Cockroaches and Robots. Integrative and Comparative Biology. 2008 02;

Lockey JK, Willis MA. One antenna, two antennae, big antennae, small: total antennae length, not bilateral symmetry, predicts odor-tracking performance in the American cockroach *Periplaneta americana*. J Exp Biol. 2015 Jul;

Louis M, Huber T, Benton R, Sakmar TP, Vosshall LB. Bilateral olfactory sensory input enhances chemotaxis behavior. Nat Neurosci. 2008 Feb;

MacIver MA, Shirgaonkar AA, Patankar NA. Energy-information trade-offs between movement and sensing. PLoS Computational Biology.

Madsen P, Johnson M, De Soto NA, Zimmer W, Tyack P. Biosonar performance of foraging beaked whales (Mesoplodon densirostris). J Exp Biol. 2005;

Maimon G, Straw AD, Dickinson MH. Active flight increases the gain of visual motion processing in Drosophila. Nature neuroscience. 2010;

Martin H. Osmotropotaxis in the Honey-Bee. Nature. 1965;

Mathew G, Mezic I. Metrics For Ergodicity and Design of Ergodic Dynamics For Multi-agent System. Physica D-Nonlinear Phenomena. 2011 Feb;

Mavrommati A, Tzorakoleftherakis E, Abraham I, Murphey TD. Real-Time Area Coverage and Target Localization using Receding-Horizon Ergodic Exploration. IEEE Transactions on Robotics. 2018;

McInroe B, Astley HC, Gong C, Kawano SM, Schiebel PE, Rieser JM, Choset H, Blob RW, Goldman DI. Tail use improves performance on soft substrates in models of early vertebrate land locomotors. Science. 2016 Jul;

Miller LM, Silverman Y, MacIver MA, Murphey TD. Ergodic Exploration of Distributed Information. IEEE Transactions on Robotics. 2016 Feb;

Mongeau JM, Frye MA. Drosophila spatiotemporally integrates visual signals to control saccades. Current Biology. 2017;

Monosov IE, Leopold DA, Hikosaka O. Neurons in the Primate Medial Basal Forebrain Signal Combined Information about Reward Uncertainty, Value, and Punishment Anticipation. J Neurosci. 2015 May;

Najemnik J, Geisler WS. Optimal eye movement strategies in visual search. Nature. 2005;

Nelson ME, MacIver MA. Sensory acquisition in active sensing systems. Journal of Comparative Physiology A. 2006;

Niven JE, Laughlin SB. Energy limitation as a selective pressure on the evolution of sensory systems. J Exp Biol. 2008;

Nyakatura JA, Melo K, Horvat T, Karakasiliotis K, Allen VR, Andikfar A, Andrada E, Arnold P, LaustrÃ¶er J, Hutchinson JR, Fischer MS, Ijspeert AJ. Reverse-engineering the locomotion of a stem amniote. Nature. 2019;

Pluta SR, Kawasaki M. Multisensory enhancement of electromotor responses to a single moving object. J Exp Biol. 2008;

Porter J, Craven B, Khan RM, Chang SJ, Kang I, Judkewitz B, Volpe J, Settles G, Sobel N. Mechanisms of scent-tracking in humans. Nat Neurosci. 2007;

Postlethwaite CM, Psemeneki TM, Selimkhanov J, Silber M, MacIver MA. Optimal movement in the prey strikes of weakly electric fish: a case study of the interplay of body plan and movement capability. Journal of the Royal Society Interface. 2009;

Ramcharitar JU, Tan EW, Fortune ES. Effects of global electrosensory signals on motion processing in the midbrain of Eigenmannia. J Comp Physiol A Neuroethol Sens Neural Behav Physiol. 2005;

Ramdya P, Thandiackal R, Cherney R, Asselborn T, Benton R, Ijspeert AJ, Floreano D. Climbing favours the tripod gait over alternative faster insect gaits. Nat Commun. 2017 Feb; Article.

Reynolds AM, Frye MA. Free-flight odor tracking in Drosophila is consistent with an optimal intermittent scale-free search. PloS one. 2007;

Robinson EL. Data analysis for scientists and engineers. Princeton University Press

Rose GJ, Canfield JG. Longitudinal tracking responses of the weakly electric fish, *Sternopygus*. J Comp Physiol A Neuroethol Sens Neural Behav Physiol. 1993;

Rucci M, Victor JD. The unsteady eye: an information-processing stage, not a bug. Trends Neurosci. 2015 Apr;

Schnitzler HU, Moss CF, Denzinger A. From spatial orientation to food acquisition in echolocating bats. Trends Ecol Evol. 2003;

Scott SE. Different perspectives and formulas for capturing deviation from ergodicity. SIAM Journal on Applied Dynamical Systems. 2013;

Sefati S, Neveln ID, Roth E, Mitchell TR, Snyder JB, MacIver MA, Fortune ES, Cowan NJ. Mutually opposing forces during locomotion can eliminate the tradeoff between maneuverability and stability. Proc Natl Acad Sci USA. 2013 Nov;

Sengupta B, Stemmler M, Laughlin SB, Niven JE. Action potential energy efficiency varies among neuron types in vertebrates and invertebrates. PLoS Comput Biol. 2010 Jul;

Shannon CE, Weaver W. The mathematical theory of communication. Urbana,: University of Illinois Press

Sponberg S, Dyhr JP, Hall RW, Daniel TL. Luminance-dependent visual processing enables moth flight in low light. Science. 2015 Jun;

Sprayberry JDH, Daniel TL. Flower tracking in hawkmoths: behavior and energetics. J Exp Biol. 2007;

Srinivasan M, Ruina A. Computer optimization of a minimal biped model discovers walking and running. Nature. 2006;

Stachenfeld KL, Botvinick MM, Gershman SJ. The hippocampus as a predictive map. Nat Neurosci. 2017 Oct; Article.

Stamper SA, Roth E, Cowan NJ, Fortune ES. Active sensing via movement shapes spatiotemporal patterns of sensory feedback. Journal of Experimental Biology. 2012 May;

Stockl AL, Kihlstrom K, Chandler S, Sponberg S. Comparative system identification of flower tracking performance in three hawkmoth species reveals adaptations for dim light vision. Philos Trans R Soc Lond, B, Biol Sci. 2017 Apr;

Thrun S, Burgard W, Fox D. Probabilistic robotics. MIT press

Vergassola M, Villermaux E, Shraiman BI. 'Infotaxis' as a strategy for searching without gradients. Nature. 2007 Jan;

Waldron KJ, Estremera J, Csonka PJ, Singh SP. Analyzing Bounding and Galloping Using Simple Models. J Mech Robot. 2008 Jul;

Watanabe A, Takeda K. The change of discharge frequency by A.C. stimulus in a weak electric fish. J Exp Biol. 1963;

Webb B, Harrison RR, Willis MA. Sensorimotor control of navigation in arthropod and artificial systems. Arthropod Struct Dev. 2004;

Willis MA, Avondet JL. Odor-modulated orientation in walking male cockroaches Periplaneta americana, and the effects of odor plumes of different structure. J Exp Biol. 2005;

Yang SCH, Lengyel M, Wolpert DM. Active sensing in the categorization of visual patterns. eLife. 2016 feb;

Yovel Y, Falk B, Moss CF, Ulanovsky N. Optimal localization by pointing off axis. Science. 2010;

Yovel Y, Falk B, Moss CF, Ulanovsky N. Active control of acoustic field-of-view in a biosonar system. PLoS Biol. 2011 Sep; PMC3172196[pmcid].

Materials and Methods

Electric fish electrosensory tracking

Three adult glass knifefish (*Eigenmannia virescens*, Valenciennes 1836, 8–15 cm in body length) were obtained from commercial vendors and housed in aquaria at ~28°C with a conductivity of ~100μS cm^{-1}. All experimental procedures were approved by the Institutional Animal Care and Use Committee of Northwestern University.

An experimental setup was built (similar to a prior study (***Stamper et al., 2012***)) in which a 1D robot-controlled platform attached to a refuge allows precise movement of the refuge under external computer control. Three adult glass knifefish (*Eigenmannia virescens*, Valenciennes 1836) were used, 8–15 cm in body length. The refuge was a customized rectangular section, made by removing the bottom surface of a 15 cm long by 4.5 cm high by 5 cm wide PVC section (3 mm thick) and making a series of 6 openings (0.6 cm in width and spaced 2.0 cm apart) on each side. These windows provide a conductive (water) alternating with non-conductive (PVC) grating to aid electrolocation. The bottom of the refuge was 0.5 cm away from the bottom of the tank to ensure that the fish stays within the refuge. A high-speed digital camera (FASTCAM 1024 PCI, Photron, San Diego, USA) with a Nikon 50 mm f/1.2 fixed focal length lens was used to capture video from below the tank viewing up at the underside of the fish (Fig. 3A). Video was recorded at 60 frames s^{-1} with a resolution of 1024 × 256 pixels. The refuge was attached to a linear slide (GL20-S-40-1250Lm, THK Company LTD, Schaumburg, USA), with a 1.25 m ball screw stroke and a pitch of 40 mm per revolution. The slide is powered by an AC servomotor (SGM-02B312, Yaskawa Electric Corporation, Japan) and servomotor controller (SGD-02BS, Yaskawa Electric Corporation, Japan). The refuge trajectory was controlled by a remote MATLAB xPC target with a customized Simulink model (MathWorks, Natick, USA).

Fish were housed in aquaria at ~28°C with a conductivity of ~100μS cm^{-1}. Before each session, individual fish were placed into an experiment tank (with identical water conditions) equipped with the moving refuge, high-speed camera and closed-loop jamming system (see below), and allowed to 2–4 hours to acclimate. Trials were done in the dark with infrared LEDs (λ = 850 nm) used to provide illumination for the camera. Each trial was 80 seconds long with the jamming signal only applied after the first 10 seconds and removed for the final 10 seconds. A total of 21 trials ($n = 10$ for strong signal and $n = 11$ for weak signal) were used for this analysis. During each trial, the servomotor drives the refuge in a predefined 0.1 Hz sinusoidal fore-aft motion with an amplitude of 17 mm.

Video of electric fish refuge tracking was processed by a custom machine vision system written in MATLAB to obtain the fish head centroid and location of the refuge at 60 Hz. The x (longitudinal) position of the centroid of head of the fish was filtered by a digital zero-phase low-pass IIR filter with a cut-off frequency of 2.1 Hz and then aligned with the refuge trajectory. For all the completed trials ($n = 21$) across a total of 3 electric fish, the trajectory of both the fish and the refuge trajectory was used for the frequency domain analysis analysis (Fig. 4). We used the Fourier transform to analyze the fish's tracking response in the frequency domain. Trials with no jamming are categorized as the strong signal condition ($n = 10$, average trial duration 59.6 seconds). Trials with jamming (jamming amplitude ≥ 10 mA, see below) are categorized as the weak signal condition ($n = 11$, average trial duration 54.5 seconds). The cumulative distance traveled by the fish and refuge during refuge tracking was computed and denoted by D_{f} and D_{r}, respectively. Relative exploration was then defined as $D_{\mathrm{f}}/D_{\mathrm{r}}$.

Closed-loop jamming system

During refuge tracking, the fish's electric organ discharge (EOD) signal was picked up by two bronze recording electrodes and amplified through an analog signal amplifier (A-M Systems Inc, Carlsborg, USA) with a linear gain of 1000 and a passband frequency from 100 Hz to 1000 Hz. A data acquisition unit (USB 6363, National Instruments, Austin TX, USA) provided digitized signals used in a

custom MATLAB script (MathWorks, Natick MA, USA) to identify the principal frequency component of the EOD. A sinusoidal jamming signal was then generated through the same USB interface using the digital to analog voltage output channel. The jamming signal's frequency was set to be a constant 5 Hz below the fish's EOD frequency as previously found most effective (***Bastian, 1987***; ***Ramcharitar et al., 2005***). The synthesized jamming signal was sent to a stimulus isolator (A-M Systems Inc, Carlsborg WA, USA), which also converted the voltage waveform into a current waveform sent to two carbon electrodes aligned perpendicular to the EOD recording electrodes to avoid interference. The efficacy of the jamming stimulus was verified by examining its affect on the EOD frequency shown in Fig. S10.

Mole olfactory localization

Tracking data of blind eastern American moles (*Scalopus aquaticus*, Linnaeus 1758) locating a stationary odor source were digitized from a prior study (***Catania, 2013***). Three experimental conditions were used in the original study: one in which there was normal airflow (categorized in the strong signal condition), one where one nares was blocked (weak signal condition), and one where the airflow was crossed to the nares using an external manifold (also weak signal condition). Relative exploration was defined as the ratio between the cumulative 2D distance traveled by mole and a straight line from its starting position to the odor location $D_{\text{mole}}/D_{\text{line}}$, to allow direct comparison between strong and weak signal conditions despite differences in the mole's initial position and target location across trials (Fig. 3C). For the corresponding EIH simulation trials, we define relative exploration as the raw cumulative lateral distance traveled by the sensor since the simulation is done in 1D.

Cockroach olfactory localization

American cockroach (*Periplaneta americana*, Linnaeus 1758) odor source localization behavior data was acquired from (***Lockey and Willis, 2015***). The cockroach's head position was tracked during an odor source localization task. The same behavioral experiments were conducted with the odor sensory organ, the antennae, bilaterally cut to a specified length. The control group with intact antennae (4 cm in length) was categorized as the strong signal condition, and the 1 cm and 2 cm antenna-clipped groups were categorized as the weak signal condition. Only successful trials ($n = 51$, 20 strong signal condition trials and 31 weak signal condition trials) were included in the analysis. Relative exploration for the cockroach data shown in Fig. 3C is computed as the ratio of the cockroach's total walking distance and the reference path length (the shortest path length from position at the start to the target) $D_{\text{cockroach}}/D_{\text{reference}}$ reported from (***Lockey and Willis, 2015***).

Hawkmoth flower tracking

Hummingbird hawkmoths (*Macroglossum stellatarum*, Linnaeus 1758) naturally track moving flowers in the wind as they insert their proboscis in the nectary to feed, primarily driven by vision and mechanoreception (***Sponberg et al., 2015***; ***Stockl et al., 2017***). In a prior study (***Stockl et al., 2017***), the hawkmoth's flower tracking behavior was measured under various levels of ambient illumination while it fed from a nectary in a robotic flower. The robotic flower moved in a predefined sum-of-sine trajectory composed of 20 prime multiple harmonic frequencies from 0.2 to 20 Hz. The moth's lateral position was tracked and the Bode gain of the raw tracking data was acquired from (***Stockl et al., 2017***) and used in our analysis. A segment of the moth's raw tracking trajectory is shown in Fig. 1. We classified trials under high illumination (3000 lx, $n = 13$) as the strong signal condition, and trials under low illumination (15 lx, $n = 10$) as the weak signal condition.

Simulations

The algorithmic implementation of EIH is built upon a framework we introduced in prior work for robotic tracking of stationary targets using Fisher information (***Miller et al., 2016***). The original algorithm was modified to track moving targets using entropy reduction as the information measure for better comparison to infotaxis, which also used this approach (***Vergassola et al., 2007***) (the results

are insensitive to the choice of information metric; near identical results were obtained with Fisher Information). Code to reproduce these simulations, the empirical data, and a Jupyter Notebook stepping through how the EID is calculated, is open-sourced (***Chen et al., 2019***). Supplementary Movie 1 provides a theory overview and Supplementary Movie 2 shows how EIH is applied to the control of sensing-related movements of a bio-inspired electrolocation robot. For pseudocode of EIH and simulation parameters, see Alg. S1 and Tables S1–S2.

For all analyzed animals, the body or sensory organ being considered is modeled as a point mass in a 1-dimensional workspace. The workspace is normalized to 0 to 1 for all the simulations. Each sensory measurement V is drawn from a Gaussian function that models the signal coming in to the the sensory system plus a zero-mean Gaussian measurement noise ϵ to simulate the effect of sensory noise:

$$V = \Upsilon(\theta, x) + \epsilon$$

where $\Upsilon(\theta, x)$ denotes the Gaussian observation model function evaluated at position x given the target stimulus location θ

$$\Upsilon(\theta, x) = \frac{1}{\sigma_{\mathrm{m}}\sqrt{2\pi}} \exp\left(-\frac{(x-\theta)^2}{2\sigma_{\mathrm{m}}^2}\right)$$

where σ_{m} is the variance of the observation model. If we fix the target location θ at location 0.5 (center of the normalized workspace) and vary x to take continuous sensory measurements from 0 to 1, the resulting measurement versus location x will form a Gaussian (Fig. 2C, upper inset).

The variance σ_{n} of ϵ is controlled by the signal-to-noise ratio (SNR) of the simulation:

$$\sigma_n = \gamma\sqrt{\frac{\sigma_m^2}{10^{\frac{\mathrm{SNR_{dB}}}{10}}}}$$

σ_{n} is the variance of the simulated sensory noise ϵ, and γ is a unity constant in units of normalized sensor signal unit per normalized workspace unit that converts σ_{m} (in normalized workspace units) on the RHS to normalized sensor signal units of the LHS term σ_{n}. We used $\sigma_{\mathrm{m}} = 0.06$ for all simulations (Table S1); The SNR values used for all the simulations is documented in Table S2. It should be noted that the values of SNR used in EIH simulations are only intended to relate qualitatively ("strong" or "weak" signal) to the actual (unknown) SNR of the animal's sensory system during behavior experiments. To explore the impact of our assumptions regarding SNR, Fig. S3 provides a sensitivity analysis showing how relative exploration varies as a function of SNR.

This generic Gaussian model of observation abstracts the process by which an animal relates afferent signals from sensory receptors to the estimated location of the target. The SNR of the observation model abstracts the effect of measurement uncertainty in the form of additive zero-mean Gaussian sensory noise. We used 10–30 dB as the weak signal condition and 50–70 dB as the strong signal condition due to the fact that relative exploration plateaus below 30 dB and beyond 50 dB SNR (Fig. S3A). We only intend to use the relative change between high and low SNRs to simulate similar changes in the behavioral conditions of strong and weak signal trials.

Information about each measurement is quantified in the form of expected information density and used through a Bayesian filter (***Thrun et al., 2005***) to acquire the posterior belief distribution, as was the case for the infotaxis algorithm (***Vergassola et al., 2007***). During a given tracking task, the EIH algorithm calculates an expected information density (EID) that represents the predicted information gain at every location in space (steps in calculating EID below; online interactive tutorial ***Chen et al.*** (***2019***), and Alg. S1). Using an iterative process, a trajectory is generated that minimizes the divergence between two probability distributions: the statistics of the trajectory, and the statistics of the EID.

The initial condition for all the simulation trials was a uniformly flat (uninformative) belief and an initial state of zero velocity and acceleration. To ensure uniformity, all the simulation trials were

set to have the exact same internal parameters except for SNR, which was changed across trials to compare trajectories in strong and weak signal conditions, and T, the duration of each planned trajectory, which was shortened for the moth trials due to the multiple frequency content of the prescribed target motion (Table S1).

Non-technical description of EIH

For each species, we model only one sensory system, the sensory system whose input was degraded through some experimental manipulation during the study. We model the sensory system as a point-sensor moving in one dimension (electrosense for fish, olfaction for mole and cockroach, and vision for moth). The sensory system is assumed to provide an estimate of location, modeled by drawing values from a Gaussian probability distribution with a variance determined by the specified signal-to-noise ratio (SNR). This is called the "observation model" for a system. Assuming this observation model and an initially uninformative ("flat") probability distribution of where the target is believed to be (for brevity, this distribution is called "the belief"), EIH proceeds as follows: 1) For the current belief, derive the corresponding EID by calculating the answer to the question "how much information (quantified by entropy reduction) can we obtain by taking a new observation at this location?" for all possible locations; 2) run the trajectory optimization solver to generate a trajectory segment with duration T (Table S1) that optimally balances energy expenditure against ergodicity with respect to the EID; 3) execute the generated trajectory, allowing the sensor to make observations along it; 4) use the incoming observations to update the belief using a recursive Bayesian filter (***Thrun et al., 2005***). This is the step that updates where EIH believes the target to be located by taking into account new evidence and existing prior knowledge; 5) Check whether the termination condition has been reached (either running for a specified time or until the variance of belief is below a certain threshold), and if not, return to step 1. Supplementary Movie 2 shows these steps graphically for control of an electrolocation robot localizing a stationary target.

Quantifying information for the Expected Information Density (EID)

Consider the case of an animal tracking a live prey. Suppose that in open space the signal profile of the prey is similar to a 3D Gaussian centered at its location. For a predator trying to localize the prey, sampling only at the peak of the Gaussian is problematic because while the signal is strongest at that location, it is also locally flat, so small variations in the prey's location has little impact on sensory input. In contrast, any motion of the prey relative to the predator at the maximum slope of the signal profile will result in the largest possible change in the signal received by the predator, and therefore maximize the predator's localization accuracy (Fig. 2C, the expected amount of information is maximal at the peak of the spatial derivative of the Gaussian profile).

Suppose at time t one has a probability distribution that is the belief $p(x)$ about the value of x, for instance about the location of a particular chemical source, prey, or predator. In EIH, a Bayesian filter is used to optimally update $p(x)$ from measurements, so $p(x)$ evolves dynamically over time (***Kording and Wolpert, 2004***). The entropy of $p(x)$, defined by $\sum_i p(x_i)\log(p(x_i))$ (where i is an index over a discretization of the domain), is the amount of information required to describe $p(x)$ as a probability distribution. The entropy of a uniform or flat distribution is high—if it represents object location, it means an object could be at all possible locations in space and—like white snow on a TV screen—requires a lot of information to describe; while a narrow distribution for an object's location can be described with very little information. The EID can be derived by simulating a set of possible sensing locations in the workspace, and for each location predicting the expected information gain by evaluating the reduction in entropy of the posterior with respect to the current prior (Appendix 4).

Ergodicity

The ergodicity of a trajectory $s(t)$ with respect to a distribution of the informativeness of sensing locations through space $\Phi(x)$ is the property that the spatial statistics of $s(t)$—the regions the tra-

jectory visits and how often it visits them—matches the spatial distribution $\Phi(x)$. Technically, this is quantified by saying that a trajectory $s(t)$ is ergodic with respect to $\Phi(x)$ if the amount of time the trajectory spends in a neighborhood $\mathcal{N}$ is proportional to the measure of that neighborhood $\int_{\mathcal{N}} \Phi(x)dx$ (Fig. 2A). With a finite time horizon, perfect ergodicity is impossible unless one uses infinite velocity, which motivates a metric on ergodicity (***Scott, 2013***). A metric on ergodicity should be zero when a trajectory is perfectly ergodic and strictly positive and convex otherwise, providing a criterion that can be optimized to make a trajectory as ergodic as possible given the control cost constraint (see below). A standard metric used for comparing distributions is the Sobolev space norm, which can be computed by taking the spatial Fourier transform of $\Phi(x)$ and $s(t)$ (see below). This metric is equivalent to other known metrics such as those based on wavelets (***Scott, 2013***). We can generate an ergodic information harvesting trajectory by optimizing the trajectory with respect to the ergodic metric (***Miller et al., 2016***), often with real-time performance (***Mavrommati et al., 2018***), both in deterministic and stochastic settings (***De La Torre et al., 2016***). See next section for more background on ergodicity.

Balancing energy expenditure and proportional betting

In EIH, candidate trajectories are generated (step 2 in the paragraph above) by minimizing the weighted sum of 1) the ergodic metric, which quantifies how well a given trajectory does proportionally betting on the EIH; and 2) the square norm of the control effort. Note that mass is implicitly included in the weighted sum. Optimizing the ergodic metric alone forms an ill-posed implementation as this implies that energy consumption is not bounded. This is equivalent to a situation where the energetic cost of movement is zero, with a consequent movement strategy of sensing everywhere. This is unlikely to be a reasonable movement policy for animals to maximize their chance of survival. Similarly, EIH is not minimizing energy either, as the energy minimizing solution alone is to not move at all. More realistically, when animals have a limited energy budget for movement, the control cost term should be added to impose a bound for energy consumption for a given trajectory. In the first-order approximation of the kinematics of motion of animals, the control cost is defined by the total kinetic energy required to execute the input trajectory (Alg. S1). In our study, the control cost is not intended to explicitly model the energy consumption of any particular animal used in the study. It is used, however, to represent the fact that energy is a factor that animals need to trade-off with information while generating trajectories for sensory acquisition. The trade-off between ergodicity and energy of motion is represented by R/λ, where R is the weight on the control cost and λ is the weight on the distance from ergodicity (Table S1). We used a value of $R/\lambda = 2$, resulting in a relative exploration value of around 2. The variation in relative exploration with an order of magnitude change in R/λ from 1 to 10 is 1.4 to 2.3 (Fig. S6).

Behavioral trajectory simulations

It is worth noting that in EIH, the animal's tracking behavior is hypothesized to be the outcome of a dynamical system, the result of forces and masses interacting, rather than sample paths of a random process—the traditional venue for ergodicity and entropy to play a role in analysis. However, we discuss the possibility that sense organs are moved stochastically in Appendix 5. When used to simulate behavioral trajectories, EIH was reconfigured to use the prescribed stimulus path from the corresponding live animal experiment as the target trajectory (Fig. 3). The simulated sensor's initial position was set to match the animal's starting location. To simulate the effect of a weak sensory signal, the SNR was reduced in the respective trials to simulate the effect of increased measurement uncertainty. Other than target trajectory, initial position, and SNR, the simulation parameters were the same across all simulations (Table S2).

Wiggle attenuation simulations

The simulated electric fish tracking response is filtered through zero-phase IIR low-pass filters with different stop band attenuations (Fig. S7A). These filters are configured to pass the low frequency tracking band within ~0.1 Hz (target motion is a sinusoid in 0.05 Hz). This configuration allows effec-

tive removal of higher frequency wiggle motion without affecting the baseline tracking response. The effect of the wiggle filter is parameterized by the stop band attenuation at 1.5 Hz. The wiggle magnitude can be systematically deteriorated by controlling the stop band attenuation while maintaining intact baseline tracking (Fig. S7A-C).

The raw simulated weak signal tracking trajectory is first filtered by the wiggle filter at stepped attenuation levels from 5 dB to 150 dB. The filtered trajectory is then prescribed to a tracking-only simulation where the sensor is instructed to move along the predefined input trajectory, take continuous sensor measurements, and use these to update the belief and EID. The distance from ergodicity is then evaluated based on the trajectory segment and simulated EID in the same way as for the other behavior simulations. Tracking performance is evaluated by comparing the sensor's best estimate of the target's position over time based on its belief and the ground truth.

Infotaxis simulations

The original infotaxis algorithm (***Vergassola et al., 2007***) was adapted for 1D tracking simulations. The infotaxis algorithm computes the EID in the same way as EIH, but differs in the movement policy once the EID is computed. The sensor considers three movement directions from its current position—left, right, or stay—at every planning update. The sensor follows the infotaxis strategy by choosing a movement direction that will maximize the EID and then takes samples along the chosen direction to update the Bayesian filter (***Thrun et al., 2005***) and consequently the EID for the next planning iteration. The parameters of the infotaxis simulation are kept the same as for EIH to allow direct comparison.

Energetics

We analyzed how the additional movement for tracking in weak signal conditions affected energy use for electric fish (Fig. 6). We estimated the net mechanical work required to move the fish along the observed tracking trajectory by first computing the instantaneous power $P(t) = F(t)v(t)$ of tracking at every timestamp. The net force $F(t)$ was estimated by applying Newton's law $F = ma$ with the estimated body mass m (from ***Postlethwaite et al.*** (***2009***)). Finally, total mechanical work done by the fish is the integral of the instantaneous power over time $\int_t P(t)dt$. The effect of added mass was included using equations previously developed (***Postlethwaite et al., 2009***).

The relative energy was defined as the total mechanical work of moving the fish along the tracking trajectory divided by the work of moving the fish along the trajectory of the target (the refuge). A relative energy of "1x" therefore indicates that moving the fish along the tracking trajectory required the same energy as moving it along the path that the moving refuge took.

Spectral analysis

The frequency response of electric fish, mole, cockroach, and moth tracking and simulation data were analyzed using the Fourier transform. The magnitude frequency response data was used in Fig. 4A-L. For the 2-D trajectories of mole and cockroach, the lateral tracking response was analyzed separately alongside the 1-D EIH lateral tracking simulation (Fig. 4E-L). Because our simulations assume a normalized workspace dimension of 0 to 1, the spectral analyses are shown with normalized Fourier magnitudes and are only intended to provide a qualitative link between EIH and animal behavior, rather than matching the units of the original live animal trajectories. For the moth, since the sum-of-sine stimulus covers a wide frequency range that includes the frequencies of the sensing-related movements, the tracking response of moth behavior and simulation is shown in the form of a Bode gain plot (Fig. 3M-N) instead of Fourier magnitude to visualize both the frequency spectrum of motion and relative exploration of each tracked frequency component. A Bode gain plot shows the magnitude of the frequency response of the tracking trajectory normalized by the stimulus for a wide range of frequencies. A gain of 1 for any particular frequency indicates the moth (or simulated sensor) responded with the same amplitude as the sum-of-sine stimulus at that frequency. The averaged Fourier magnitude and Bode gain were computed by taking the mean of the Fourier magnitudes or Bode gain within the wiggle frequency window marked

by the shaded area of the spectrum plots shown in the first columns of Fig. 4. For electric fish, the wiggle frequency window is identified as high frequency components outside of the baseline tracking response frequency range (Fig. 4A-B). For the mole and cockroach, because the target is stationary and hence there is no baseline tracking frequency, the entire frequency spectrum of the tracking response was used for computing the statistics.

Quantification and statistical analysis

The Kruskal-Wallis one-way ANOVA test was used for the statistical analysis of relative exploration (Fig. 3) and spectral power of tracking (Fig. 4). Each trial of weakly electric fish, mole, cockroach, and moth behavior as well as their corresponding simulations were considered independent. Kruskal-Wallis is non-parametric and hence can be applied to test for the significance of relative exploration even though it is a ratio distribution.

The Pearson correlation coefficient and the 95% confidence interval of its distribution were calculated in Fig. S3B based on data from Fig. S3A. The mean and 95% confidence interval was computed for Fig. 3 and S3.

Data and software availability

All data and code needed to reproduce our results is available (***Chen et al., 2019***), as well as an interactive Jupyter notebook tutorial on computing the EID. Algorithm S1 provides psuedocode, and Tables S1-S2 provide the corresponding simulation parameters for the EIH algorithm. Finally, Supplementary Video S1 provides a video explainer of the theory of EIH, and Video S2 applies it to controlling an underwater electrolocation robot.

Appendix 1

Background on ergodicity

Ergodicity plays an important role in multiple scientific disciplines particularly in stochastic systems and statistical mechanics. In the setting of Markov chains, defined by states and stochastically-driven transitions between states, a system is ergodic if every aperiodic path that leaves a given state must return to that state with probability one (***Lasota and Mackey, 2013***). However, in the present work we are not interested in stochastic evolutions—though there is the possibility that stochasticity in a system could contribute to coverage needs, something discussed momentarily. Instead, we are interested in deterministic decisions—i.e., control decisions—that provide coverage with respect to regions of high information density.

A key insight from (***Mathew and Mezic, 2011***) was to use the definition of ergodicity for dynamical systems—that a trajectory $x(t)$ spends time in any particular neighborhood $\mathcal{N}$ proportional to the measure of the distribution Φ over that neighborhood $\int_{\mathcal{N}} \Phi(x)dx$—to create a metric on deterministic trajectories. That is, once Φ is given, there is nothing stochastic in the question of coverage. Instead, there is only the question of how much coverage a particular trajectory $x(t)$ provides relative to Φ. It should additionally be noted that coming up with a metric is not trivial, in large part because any mathematical comparison must be able to compare two distinct mathematical ideas—a distribution and a trajectory. A distribution is a probability over a region, while a trajectory is a continuum of states parameterized by time t. In (***Mathew and Mezic, 2011***) the authors note two critical steps in creating such a comparison. First, they note that a trajectory can *also* be represented as a sequence of Dirac delta functions $\delta(x)$ also parameterized by time t, so that the comparison is between two distributions rather than between a distribution and a trajectory. Secondly, perhaps more importantly, they use the fact that spatial Fourier transforms are well posed for quite general distributions, including Dirac delta functions. From these two steps they conclude that the coefficients of the Fourier transform provide an infinite set of variables that can be used to form a metric. Importantly, none of this analysis requires the trajectories to be stochastic, though similar analysis can be done for stochastic executions (where, for instance, one can imagine that

at least small amounts of noise might improve coverage locally).

Appendix 2

Relationship between Ergodicity and Kullback-Leibler divergence

Ergodicity provides a mathematical approach for comparing a trajectory to a distribution in a way that is similar to how K-L divergence compares a distribution P to a distribution Q. However, K-L divergence cannot be directly applied to trajectories. This is because an idealized trajectory (no uncertainty by itself) is an aggregations of singletons in the form of individual Dirac delta functions (each of zero variance), one for each time t. This leads to infinite K-L divergence because the K-L divergence measures how much information changes when using one distribution to represent another distribution (this is why the K-L divergence is often called the relative entropy). In the case of representing a distribution with a delta function, infinite information has been gained because the argument of the delta function is specified with zero variance. As a concrete special case, if one approximates a delta function with normal distributions of decreasing variance, the entropy goes to infinity as the variance goes to zero.

Here is a brief illustration of some of the steps needed to show why K-L divergence will not work for trajectories for those who would like more explanation. This illustration elides a number of technical issues that would need to be carefully worked through for a more rigorous result. Imagine we have a 1D trajectory that consists of two points (*i.e.* singletons): $\{0, 1\}$. We can call this a probability density function (PDF) P consisting of two Dirac deltas, $\delta(x-0)$ and $\delta(x-1)$—that is, $P(x) = \frac{1}{2}(\delta(x-0) + \delta(x-1))$. The differential probabilities $P(0) = P(1)$ both integrate to $\frac{1}{2}$ since the total probability is 1. Suppose we want to compute the K-L divergence between P and Q, where Q is an arbitrary Gaussian distribution with a mean of 0.5 and a non-zero variance. According to the general form of computing K-L divergence:

$$\begin{aligned} D_{KL}(P||Q) &= \int_x P(x) \log \frac{P(x)}{Q(x)} dx \\ &= \int_x P(x) \log P(x) dx - \int_x P(x) \log Q(x) dx \end{aligned} \tag{1}$$

where $\int_x P(x) \log P(x) dx$ is the differential entropy of P, which is undefined in this case. To understand why, consider an arbitrary Gaussian distribution, $f(x) \sim \mathcal{N}(\mu, \sigma)$. Computing the first term of the expression for differential entropy (Eq. 1) gives $h(f) = \int_x f(x) \log f(x) dx = \log(\sigma) + \frac{1}{2}\log(2\pi e)$, which is undefined in the limiting case of a Dirac delta function with $\sigma = 0$ since $\lim_{\sigma \to 0^+} \log(\sigma) = -\infty$. Hence, the K-L divergence between a Dirac delta function (representing the idealized trajectory) and a smooth (EID) distribution is undefined. (Note that the other term in the K-L divergence that depends on $Q(x)$ will evaluate to a constant, so does not impact the well-posedness of K-L divergence for a trajectory.)

Similar to K-L divergence, mutual information, which quantifies the amount of information obtained about one random variable X by observing another random variable Y, defined as $I(X;Y) = D_{\mathrm{KL}}\left(P(X,Y) \,||\, P_X P_Y\right)$, is another widely used approach for quantifying information between two distributions of random variables. For jointly discrete or jointly continuous pairs (X, Y), it is the K-L divergence between the joint distribution $P(X, Y)$ and the product of the marginal distributions P_X and P_Y. Given that the K-L divergence between a trajectory and a distribution is undefined as discussed above, mutual information also cannot be applied to trajectories. More generally, as we stated in the main text that animal physical trajectories are considered the behavior of dynamical systems rather than sample paths from a stochastic process, methods like K-L divergence and mutual information that requires both inputs to be distributions are undefined and hence will not work in the case where one of these is a trajectory.

Appendix 3

How the distance from ergodicity is computed

These details are largely from our prior publication ***Miller et al.*** (***2016***), repeated here for convenience. The spatial statistics of a trajectory $x(t)$ are quantified by the percentage of time spent in each region of the workspace

$$C(x) = \frac{1}{T}\int_0^T \delta[x - x(t)]dt \tag{2}$$

where δ is the Dirac delta function and T is the duration of the trajectory. The distance from ergodicity is then defined as the sum of weighted squared distances between the Fourier coefficients of the EID ϕ_k and the coefficients c_k of the distribution representing the time-averaged trajectory. The Fourier coefficients $\phi_{\boldsymbol{k}}$ of the distribution $\Phi(\boldsymbol{x})$ are computed using an inner product

$$\phi_{\boldsymbol{k}} = \int_X \phi(\boldsymbol{x}) F_{\boldsymbol{k}}(\boldsymbol{x}) d\boldsymbol{x} \tag{3}$$

and the Fourier coefficients of the basis functions along a trajectory $\boldsymbol{x}(t)$, averaged over time, are calculated as

$$c_{\boldsymbol{k}}(\boldsymbol{x}(t)) = \frac{1}{T}\int_0^T F_{\boldsymbol{k}}(\boldsymbol{x}(t))dt \tag{4}$$

where T is the final time and F_k is a Fourier basis function that takes the form of

$$F_k(x) = \frac{1}{h_k}\prod_{i=1}^{n}\cos\left(\frac{k_i\pi}{L_i}x_i\right) \tag{5}$$

where h_k is a normalization factor ***Mathew and Mezic*** (***2011***) and L_i is a measure of the length of the dimension. Finally, the ergodic metric is specified as

$$\mathcal{E}(x(t)) = \sum_{k=0}^{\mathbf{K}} \Lambda_k [c_k(x(t)) - \phi_k]^2 \tag{6}$$

where $\mathbf{K}$ is the number of Fourier coefficients used along every one of the n dimensions and $\Lambda_k = (1 + \|k\|^2)^{-s}$, where Λ_k come from the Sobolev space norm and $s = \frac{n+1}{2}$ places more weight on lower frequency information ***Mathew and Mezic*** (***2011***). Given the definition above, the ergodic metric $\mathcal{E}(x(t))$ quantifies the difference between a given distribution of EID and the spatial statistics $C(x)$ of an trajectory x. We say a trajectory is perfectly ergodic with respect to an EID map if $\mathcal{E}(x) = 0$, that is, the spatial statistics of x exactly matches the distribution of the target EID map.

Appendix 4

How the expected information density (EID) is defined and computed

Given an unknown random variable $\boldsymbol{\theta}$ to estimate (in the context of tracking simulations, $\boldsymbol{\theta}$ represents the location of the tracking target), EIH evaluates an expected information density EID(x) at every planning update based on the current belief $p(\boldsymbol{\theta})$. The EID essentially answers the following question: given the probability of $\boldsymbol{\theta}$ being a particular value, and given the likelihood of receiving a particular voltage V corresponding to that value, what is the average amount of information we expect to receive by visiting a state x?

Computing EID(x) requires several steps. First, we define a Gaussian likelihood function that predicts how likely the sensor is to obtain a measurement $V \in \mathcal{V}$ given the current belief $p(\boldsymbol{\theta})$, where $\mathcal{V}$ is the set of all possible sensor measurements (see Chapter 7.2 of ***Robinson*** (***2016***) for details regarding the likelihood function).

$$p(V \mid \boldsymbol{\theta}, x) = \frac{1}{\sqrt{2\pi}\sigma}\exp\left[-\frac{(V - \Upsilon(\boldsymbol{\theta}, x))^2}{2\sigma^2}\right] \tag{7}$$

Here $x \in \mathcal{X}$ is the location of the sensor ($\mathcal{X}$ is the space of all possible sensor locations), and $\Upsilon(\boldsymbol{\theta}, x)$ is the observation model assuming a known target location $\boldsymbol{\theta}$ evaluated at sensor location x.

Next, with a predicted distribution of measurements for each choice of x from the likelihood function $p(V|\boldsymbol{\theta}, x)$, we evaluate what the new posterior belief $p(\boldsymbol{\theta} \mid V, x)$ is expected to be if the sensor were to take a hypothetical measurement at a given location x in the workspace. From the multiplication rule of conditional probability (see Eq. 3.16 of ***Kokoska and Zwillinger*** (***2000***)),

$$P(A \cap B) = P(A|B)\ P(B) = P(B|A)\ P(A) \tag{8}$$

where A and B are two random events, we obtain the Bayes update rule:

$$P(A|B) = \frac{P(B|A)\ P(A)}{P(B)}. \tag{9}$$

For each choice of potential x where a sensor measurement could be taken, the new posterior is therefore computed by (see Chapter 3.3.9 of ***Kokoska and Zwillinger*** (***2000***) and Chapter 2.4 of ***Thrun et al.*** (***2005***)):

$$\begin{aligned} p(\boldsymbol{\theta} \mid V, x) &= \eta\ p(\boldsymbol{\theta} \mid x)\ p(V \mid \boldsymbol{\theta}, x) \\ &= \eta\ p(\boldsymbol{\theta})\ p(V \mid \boldsymbol{\theta}, x) \end{aligned} \tag{10}$$

where $\boldsymbol{\theta}$ corresponds to A and V corresponds to B in Eq. 8. In Eq. 10, $p(\boldsymbol{\theta} \mid x) = p(\boldsymbol{\theta})$ because $\boldsymbol{\theta}$ and x are mutually independent, and $\eta = \frac{1}{p(V \mid x)} = \frac{1}{\int_{\boldsymbol{\theta}} p(V \mid \boldsymbol{\theta}, x)\ p(\boldsymbol{\theta}) d\boldsymbol{\theta}}$ is a normalization factor that constrains the posterior belief $p(\boldsymbol{\theta} \mid V, x)$ to be a probability distribution (see Chapter 2.4 of ***Thrun et al.*** (***2005***)).

Given a posterior belief $p(\boldsymbol{\theta} \mid V, x)$ evaluated on a potential V measured at a potential location x, the entropy reduction from the prior belief $p(\boldsymbol{\theta})$ can be evaluated using:

$$\Delta S(V, x) = S\left[p(\boldsymbol{\theta})\right] - S\left[p(\boldsymbol{\theta} \mid V, x)\right] \tag{11}$$

where $S\left[p(\boldsymbol{\theta})\right] \in \mathbb{R}^1$ and $S\left[p(\boldsymbol{\theta})\right] = -\sum_{\boldsymbol{\theta}} p(\boldsymbol{\theta})\ \log p(\boldsymbol{\theta})$ is the Shannon-Weaver entropy of the prior belief $p(\boldsymbol{\theta})$, while $S\left[p(\boldsymbol{\theta} \mid V, x)\right]$ is the Shannon-Weaver entropy of the posterior belief.

For any given prior belief $p(\boldsymbol{\theta})$, the probability of the sensor receiving a measurement V given a choice of sensing location x is not necessarily constant. Therefore, to evaluate the expected entropy reduction at a given sensing location x, the entropy reduction $\Delta S(V, x)$ needs to be weighted by the measurement probability $p(V \mid x)$ that is consistent with the prior belief $p(\boldsymbol{\theta})$. This weighted probability can be obtained by applying the law of total probability (see Eq. 3.17 of ***Kokoska and Zwillinger*** (***2000***)) to the normalized likelihood function $p(V \mid \boldsymbol{\theta}, x)$ treated as a probability distribution (see Chapter 5.2 of ***Robinson*** (***2016***)).

$$p(V \mid x) = \int_{\boldsymbol{\theta}} p(\boldsymbol{\theta}) p(V \mid \boldsymbol{\theta}, x) d\boldsymbol{\theta} \tag{12}$$

Finally, the expected information density at location x—EID(x)—is obtained by computing the mathematical expectation (see Chapter 3.5.1 of ***Kokoska and Zwillinger*** (***2000***)) of the entropy reduction *if one were* to take a measurement at location x. That is, EID(x) is the weighted average entropy reduction resulting from the conditional probability $p(\boldsymbol{\theta} \mid V, x)$, weighted by the measurement probability $p(V \mid x)$.

$$\text{EID}(x) = \text{E}\left[\Delta S(x)\right] = \int_V p(V \mid x)\ \Delta S(V, x) dV \tag{13}$$

An interactive online tutorial and video describing the above steps of EIH is available (***Chen et al., 2019***).

Appendix 5

Comparison to stochastic models

If animal trajectories are sample paths of a process made up of deterministic and stochastic parts, then some observed small-amplitude oscillations can be modeled by a stochastic search process, similar to that reported in *Drosophila* ***Reynolds and Frye*** (***2007***); ***Censi et al.*** (***2013***); ***Mongeau and Frye*** (***2017***); ***Ferris et al.*** (***2018***). Since a sensed target is always present in the tracking behavior analyzed here, it is unlikely that the trajectories analyzed here are purely driven by stochastic search with no intended target (*e.g.* refuge, food source, odor plume). Most models that implement stochasticity by drawing actions based on the EID distribution represent stochasticity in an abstraction of the space in which the body evolves based on its physics (e.g., a Thompson sampling process that randomly samples locations, ignoring the physics and energetics of getting to those locations). If a stochastic signal is directly driving the physics of the body, small random walks will indeed occur, but large-scale motion of the entire body will not occur unless the physical randomness is very large. Moreover, stochastic search can be considered a special case of ergodic search. In general, a random walk will lead to coverage of some area, and that same area could be covered using the ergodic coverage algorithms described here. But the ergodic coverage algorithms enable an animal to adapt as the environment changes, where a given stochastic search will be independent of changes in the environment. This difference may matter in settings where changes in environment matter to search success. Such a scenario is demonstrated in Fig. S9 where we provide two examples of target loss. In both cases, the animal exhibits immediate local-to-global search transitions which is naturally reproduced by EIH. Finally, the ergodic harvesting strategy can be applied to stochastic scenarios, where the dynamics include a stochastic process, as shown in ***De La Torre et al.*** (***2016***) without substantially impacting the solution.

Appendix 6

EIH in higher dimensional workspaces

The presented animal behavior analysis and EIH simulations are limited to a 1-D workspace in the presence of a single target. Although we show evidence suggesting that EIH naturally balances exploration-versus-exploitation trade-off in the case of signal loss in 1-D (see Supplementary Fig. S9), it is unclear how EIH would behave in similar cases in workspaces with more than one dimension, and whether the prediction is supported by animal behavior. For example, Calhoun et al. ***Calhoun et al.*** (***2014***) show that infotaxis will respond to a local distractor in the EID by first going straight towards the peak, as also shown in Fig. 2B, and then engage in circular motion around the distractor peak as it gradually gets rejected by new observations. As a comparison, in Fig. 2 we show that EIH predicts the animal will not wander around the distractor peak for long, but rather dwell in such distractor peaks to make additional observations before naturally switching to other regions including the true target location. Such behavior emerges from EIH without any dependence on changes in the EID, whereas infotaxis is dependent on changes in the EID. This also applies to Fig. S3 at 20 dB SNR, one of few places where infotaxis exhibits a relative exploration level similar to EIH. This increase in sensor movement is driven by changes in the information landscape due to a very high level of uncertainty. However, the movements of infotaxis at this SNR does not generate systematic coverage with respect to the EID and actually leads to sub-optimal tracking performance (64 ± 3% tracking error for EIH versus 74 ± 4% tracking error for infotaxis at 20 dB; Kruskal-Wallis test, $p < 0.001$, $n = 20$). Further investigation is needed to explore the effect of temporal change in the information landscape on sensing behavior (especially in higher dimension workspaces, *e.g.*, Calhoun et al. ***Calhoun et al.*** (***2014***) show behavior in 2-D) for insight into on how animals approach the exploration-versus-exploitation dilemma in various scenarios such as signal loss.

Titles and captions for supplementary items

Algorithm S1 Animal Tracking Simulation with Ergodic Information Harvesting

Table S1. Parameters of EIH Simulation

Table S2. Simulation Parameters Used For Each Figure

Figure S1. Illustration of differences between EIH and entropy minimization. (**A-B**) An illustration of an animal localizing a target of interest through moving its sensory organs or whole body with the proportional-betting (**A**) and information maximization (entropy minimization) (**B**) strategy. Black dots along the red trajectory indicate fixed time intervals, so longer distances between dots indicates higher speed. Uknown to the animal, the true target is the white star on the left, while on the right is a distractor (at center of red zone). The heat map represents the expected information density. Each step of EIH (Algorithm S1, **A**) is taken so that the resulting trajectory spends time in areas (*e.g.* $\mathcal{N}$) proportional to the expected information of that area. Trajectories generated by EIH, such as the one shown here, are ergodic with respect to the expected information because the percentage of time the trajectory spends in any subset $\mathcal{N}$ from $t = 0$ to $t = T$ is equal to the measure of expected information within $\mathcal{N}$; this condition must hold for all possible subsets. Information maximization (**B**), by contrast, locally maximizes the expected information (*i.e.* EID) at every step, leading it to go straight to the nearby distractor by following the gradient in this case. (**C-D**) Simulated animal tracking a target moving in 1D sine wave. Here we show simulated entropy minimization tracking trajectories alongside the corresponding simulated EIH trajectories under extreme SNR conditions (SNR of 10 dB and 55 dB, also shown in Fig. S3A). (**C**) With extremely high uncertainty (SNR $= 10$ dB), both algorithms were unable to accumulate enough information to successfully track the target. EIH results in persistent activation of movement even when the expected informativeness of sensing locations is flat due to a lack of information. The trajectory is prescribed to visit all locations of space with equal probability, covering the majority of the workspace (but not the entire workspace due to the energy constraint in trajectory optimization; see Algorithm S1). Entropy minimization results in cessation of motion under the same conditions. (**D**) Same as (C), but under extremely low uncertainty (strong signal) condition (SNR $= 60$ dB). Now entropy minimization hugs the one peak of the EID (there are two, as shown in Fig. 2C), while the EIH trajectory samples both peaks of the EID with excursions away from these peaks, similar to electric fish refuge tracking in the strong signal condition (Fig. 3A), and what can be inferred from a prior behavioral study in bat echolocation ***Yovel et al.*** (***2010***).

Figure S2. Evolution of belief over time for the trials shown in Fig. 3. EIH simulations of tracking behavior of weakly electric fish, mole, and cockroach. The trials shown are identical to those shown in Fig. 3. Simulated sensor position over time is the solid green line for the strong signal condition, and the solid red line in the weak signal condition. The target position is the dashed blue line. The repeated dashed gray lines mark each planning update event that segment the trajectory into individual planned trajectory segments of length T. The magenta heatmap shows the progression of the belief over time. Note the presence of fictive distractors, where the belief distribution becomes multi-peaked (asterisks). The inset in the mole and cockroach trials show a magnified segment to better visualize these distractors in the strong signal case.

Figure S3. Systematic comparison between EIH and infotaxis in tracking a target moving in sinusoid. Simulations of sensor tracks a target moving in sinusoid under 17 different SNR conditions from 10 dB to 55 dB. For each SNR condition, 10 simulations with different pseudo-number seeds are performed to establish the confidence interval ($n = 170$ for EIH, $n = 170$ for infotaxis). (**A**) As the SNR decreases from 55 dB, EIH exhibits elevated relative exploration (see Methods) in the form of increased wiggle amplitude. This is consistent with the animal behaviors summarized in Figures 2-3. EIH's elevated exploration with decreasing SNR tapers at ≈20 dB as the Bayesian filter fails to converge the posterior due to significantly less informative measurements. Infotaxis does not show the trend of increased exploration as signal weakens between 55–20 dB, and prescribed cessation of movement at the lowest (10 dB) signal strength. Throughout most of the frequency spectrum, infotaxis attempts to hug one edge of the EID, leading to near unity relative exploration. (**B**) The correlation coefficient between relative exploration and signal strength computed from **A** for EIH and infotaxis (*n*=170 trials for EIH, and *n*=170 for infotaxis). The vertical line indicates the 95% confidence interval. Ergodic harvesting has a significantly more negative correlation, indicating exploration increases as signal strength declines. The same trend is absent in infotaxis.

Figure S4. Dual target tracking simulation with EIH and infotaxis. In this simulation, two identical targets (with the same observation model) are present in the workspace, indicated by the two blue lines. To help visualize the outcome, both the belief (left panel) and EID (right panel) distribution over time are shown in magenta. In both cases, we use identical SNR levels corresponding to a weak signal condition (30 dB, see Table S2). (**A**) EIH tracks both targets with an oscillatory motion. In addition, belief peaks corresponding to the top and bottom targets are clearly visible in the left panel. (**B**) Infotaxis gets trapped at one of the information maximizing peaks, as seen in the right panel. This leads to cessation of movement and failure to detect the other target on the top (left panel).

Figure S5. Single target tracking simulation with EIH and infotaxis in the presence of a simulated physical distractor. In this simulation, a single target and a physical distractor coexist in the workspace. The simulated physical distractor has a different observation model that leads to a different measurement profile when compared to the desired target. To help visualize the outcome, both the belief (left panel) and EID (right panel) distribution over time are shown in magenta. In both cases, we use identical SNR levels corresponding to a weak signal condition (30 dB, see Table S2). (**A**) EIH is able to find the desired target (the continuous peak along the desired target location in the left panel) and rejects the distractor (no continuous peak along the distractor location in the left panel). (**B**) Infotaxis gets trapped at one of the information maximizing peaks, as seen in the right panel. Although it does find the desired target, it fails to reject the distractor, as seen in the continuous peak near the distractor in the left panel.

Figure S6. Sensitivity analysis on the ratio between control cost and ergodic cost in the objective function of trajectory optimization. Simulations are conducted in the same way as Fig. S3 but only for a fixed SNR under weak signal conditions (20 dB). The control cost term R is varied while fixing the ergodic cost term λ to be 5 for all the simulations. We simulated 16 different ratios of R/λ (shown along x axis on a log scale), each done with 10 different pseudo-random number seeds to establish confidence intervals. The relative exploration of the simulated trajectories are shown in the y axis on a linear scale with the sample mean marked by the solid color line and 95% confidence interval marked by the vertical color patch. The R/λ ratio used for all the EIH simulations included in the results is shown by the red dashed vertical line. Overall, as the R/λ drops, relative exploration drops monotonically because movement incurs higher cost. The trend also suggests that as the R/λ ratio increases, movement will cease. The value of $R/\lambda = 2$ that we used in our study results in a relative exploration value of around 2. The variation in relative exploration with an order of magnitude change in R/λ from 1 to 10 is 1.4 to 2.3. The same sensitivity analysis was also done for a strong signal condition but is not shown as the trend is similar.

Figure S7. How sensor wiggles were attenuated for analyzing the impact of their removal. (**A**) Response of three types of wiggle attenuation filters. We used an IIR lowpass filter with a cutoff frequency of 0.1 Hz to avoid filtering the baseline tracking frequency band of the target (≈0.05 Hz in this case). In the two filtered examples shown, the attenuation at 1.5 Hz is set to 50 dB and 150 dB to achieve progressively stronger attenuation of sensor wiggling. (**B**) Frequency response of the original and filtered trajectory. The baseline tracking gain between 0 Hz to 0.1 Hz is well preserved while the wiggling frequency band between 0.1 Hz to 1.5 Hz is attenuated according to the attenuation gain. (**C**) Comparison of original and filtered trajectories. As the stopband attenuation increases, wiggle amplitude decreases.

Figure S8. Rat odor tracking behavior and EIH simulation. In a prior study (***Khan et al., 2012***), Wistar rats (*Rattus norvegicus*, Berkenhout 1769) performed an odor tracking task by following a uniform odor trail on a moving treadmill with only olfactory information, under two different surgical conditions: with one side of the nose blocked (single-side nares stitching) and no blocking (sham stitching) as a control group. The experiment setup and tracked trajectory is shown in the top panel. Position data on the axis parallel to the treadmill were ignored, while positions on the perpendicular axis were recorded. The relative exploration was computed as with electric fish, using the cumulative 1D distance traveled by the rat divided by the cumulative distance of the transverse location of the odor trail. The olfactory system of the rat was modeled in the same fashion as for electric fish: as eventually providing an estimate of the location of the odor trail, modeled as a single 1D point-sensor of location with a Gaussian probability distribution controlled by the SNR of simulation (Methods). The blocked nares condition is here counted as the weak signal with lower SNR, and the sham stitching condition is counted as the strong signal. In the second row, we show EIH simulations of the weak and strong target trajectories, with the EID profile overlaid in magenta. It shows similar predicted sensor trajectories with a similar change in the relative exploration between these signal conditions as measured. These data are not included in the main study due to the low number of trials.

Figure S9. Measurements compared to prediction of EIH in two conditions where an animal needs to find the signal during tracking behavior. As a demonstration of how our model seamlessly transitions between exploitation and exploration, we examined an instance in the measured behavior of the live animals in which the rat appears to lose a signal it is following ***Khan et al.*** (***2012***), and an instance in which the mole starts off without finding the path to the source ***Catania*** (***2013***). The measured result shown here for both animals was wider casting motions transverse to the direction of movement. We examined whether these two behaviors (wider casting with loss of signal during tracking or when not finding the signal initially) would be predicted by EIH. (**A**) To simulate the condition of losing a signal trail in the middle of a successful tracking trial, we ran EIH as for the rat olfactory tracking case of Fig. S8, but with the trajectory shown in in panel A. In the "target lost" region (green square in figure), the sensor's input was replaced by random draws from the observation model. (**B**) To simulate the condition of starting without knowing where the trail is for an extended period of time, we ran EIH as for the mole olfactory tracking case of Fig. 3B, but with the trajectory shown in panel B. In the "target not yet acquired" region (green square in figure), the sensor's input was replaced by random draws from the observation model. In both instances, we observed wider casting motions transverse to the direction of travel similar to those measured (output from EIH below experimental data, overlaid with the EID profile in magenta). Similar to the measured behavior, our algorithm seamlessly transitions from exploitation to exploration when the expected information density becomes more diffuse.

Figure S10. Effect of jamming and how it varies with jamming intensity. (**A**) The fish's electric organ discharge (EOD) frequency shifts up continuously as the jamming signal is being applied. The area shaded with light blue indicates when jamming is turned on and lasts 40 seconds. The jamming signal frequency (current amplitude = 25 mA) is set to be 5 Hz below the fish's EOD frequency. The fish's EOD frequency gradually declines back to its normal value after jamming is turned off after around 60 seconds. (**B**) Maximum frequency shift during the fixed 40 second jamming window as a function of jamming current amplitude. Maximum frequency shift is averaged for every jamming amplitude across a total of 26 trials ($n = 6$ for 0 mA, $n = 5$ for 5 mA, $n = 3$ for 10 mA, $n = 5$ for 25 mA, $n = 7$ for 50 mA). The amount of jamming applied has a clear positive correlation with the EOD frequency shift.

Algorithm S1 Animal Tracking Simulation with Ergodic Information Harvesting

function TrialSim($\boldsymbol{\theta}(t)$, $\Upsilon(\boldsymbol{\theta}, x)$, σ_{m}, δ_t, T, α_s, β_s, λ, $\boldsymbol{R}$, d_{S}, $\alpha(0)$, SNR)

- Argument list: ground truth target trajectory $\boldsymbol{\theta}(t)$, observation model $\Upsilon(\boldsymbol{\theta}, x)$, variance of the observation model σ_{m}, time step δ_t, length of the planned trajectory segment T, step size control of the line search in the trajectory optimization α_s and β_s, weight for the ergodic metric term in the objective function λ, weight for the control cost term in the cost function $\boldsymbol{R}$, number of dimensions used for Sobolev space norm in ergodic metric d_{S}, initial control $\alpha(0)$, and SNR of the simulation SNR. See Table 1 for the value used for each of the parameters, and Table 2 for the SNR value used for each figure.
- Note: For trial simulations, the target location as a function of time ($\boldsymbol{\theta}(t)$) is set to what it was in the original experimental data set; however, the simulated animal does not know $\boldsymbol{\theta}(t)$). Where EIH is being used in the real world, within a robot or instantiated in biology, $\boldsymbol{\theta}(t)$ would not be specified.
- The observation model is $\Upsilon(\boldsymbol{\theta}, x)$, where $\boldsymbol{\theta}$ is the parameter unknown to the simulated animal (in the tracking context, target location); $x \in \mathcal{X}$ is the position of the sensor ($\mathcal{X}$ is the space of all possible sensor locations in the workspace); measurement $V = \Upsilon(\boldsymbol{\theta}(t), x) + \epsilon$ (each measurement V in simulation is synthesized based on this equation; in real world applications, this is obtained by the animal or robot sensory system), where $V \in \mathcal{V}$ and $\mathcal{V}$ is the set of all possible measurements; and ϵ is the additive zero-mean Gaussian measurement noise with variance σ_n (defined based on the SNR of the simulation and the variance of the observation model σ_{m}, see Methods for complete definition).
- Define cost function $J(x, a) = \lambda \mathcal{E}(x(t)) + \boldsymbol{R} \int_0^T \frac{1}{2} a(\tau)^2 d\tau$ for $\lambda > 0$ and $\boldsymbol{R} > 0$, where $\mathcal{E}(x(t))$ is the ergodic cost based on the current EID, and λ is the ergodic cost weighting factor (see entry for λ, Table 1), while $\boldsymbol{R} \int_0^T \frac{1}{2} a(\tau)^2 d\tau$ is the control cost weighted by $\boldsymbol{R}$ (Table 1), $a(\tau)$ is the control input at time τ that drives sensor motion
- Define initial state x_0, the maximum length of the simulation T_{max}, and T, the duration of each planned trajectory
- Define $\delta > 0$, the threshold on the norm of the gradient used to terminate the line search in the trajectory optimization procedure
- Initialize prior belief $p(\boldsymbol{\theta}_0)$ as a uniform distribution

Compute initial EID(x).

while $t \leq T_{\mathrm{max}}$ **do**

Generate an ergodic harvesting trajectory over interval $[t, t+T]$, and then update the belief and EID while executing this trajectory segment

EIH($p(\boldsymbol{\theta})$, α, T, t)

end while

end function

function EIH($p(\boldsymbol{\theta})$,α,T,t)

Argument list: belief $p(\boldsymbol{\theta})$, control input α, length of the planned trajectory T, and current time of the simulation t.

Trajectory optimization to solve for the optimal trajectory that minimizes the ergodic measure under the constraint of control cost

procedure Ergodic Trajectory Optimization

Iteratively optimize trajectories (e.g., using gradient descent and δ to terminate ***Lasdon et al.*** (***1967***)) to optimize J at the current time t with the current EID

end procedure

Execute optimized trajectory, updating the Bayesian belief at every time δt in the simulation

procedure Bayesian Update(see ***Thrun et al.*** (***2005***) for more detail on Bayesian filtering)

for all $s \in [t, t+T]$, **do**

Simulate system state $(x(s), a(s))$

Take new observation $V(s)$ (in simulation, $V(s)$ is provided to the simulated animal given its $x(s)$ from a synthesis process it does not have access to that uses $\boldsymbol{\theta}(s)$, $\Upsilon(\boldsymbol{\theta}(s), x(s))$, and ϵ; but in real world conditions, $V(s)$ is obtained from the sensory system).

Compute likelihood $p(V \mid \boldsymbol{\theta})$

Update posterior $p(\boldsymbol{\theta} \mid V)$

end for

Compute EID($x(t+T)$) for planning the next trajectory segment (see text for details).

Note: For Fig. 2, we plot the EID as it is computed at each time step for illustration purposes only (thus the EID computation is within the above loop); in either case, the updated EID only impacts the sensor trajectory after $(t+T)$ when the ergodic trajectory optimization routine is called.

end procedure

end function

Table 1. Parameters of EIH Simulation

Parameter	Symbol	Value	Source and Note
Variance of observation model	σ_m	0.06	σ_m is initially chosen to fit weakly electric fish behavior and kept the same for all the sensory modalities simulated for the sake of modal consistency
Time step of the simulation	δt	0.025, 0.005	In seconds. δt is initially chosen to fit weakly electric fish behavior and fixed for all the EIH and infotaxis simulations except for moth, where δt is set to 0.005 s to account for the increased bandwidth of the sum-of-sine trajectory
Duration of planned trajectory	T	2.5, 0.5	In seconds. T is initially chosen to fit weakly electric fish behavior and kept the same for all the EIH simulations except for moth, where T is set to 0.5 to account for the increased bandwidth of the sum-of-sine trajectory.
Step size control of the backtracking line search of trajectory optimization	α_s	0.1	α_s and β_s were picked to balance between the speed of convergence and the final cost of the trajectory optimization and are fixed across all the EIH simulations
Step size control of the backtracking line search of trajectory optimization	β_s	0.4	α_s and β_s were picked to balance between the speed of convergence and the final cost of the trajectory optimization and are fixed across all the EIH simulations
Weight of the distance from ergodicity term in the cost function of trajectory optimization loop (see Alg. S1)	λ	5	λ is initially chosen to fit weakly electric fish behavior and kept the same for all the simulations. Note that changing λ changes the trade-off between ergodicity $\mathcal{E}$ (how much information one wants) and control effort u(how much energy one is willing to give up). As a result, there is mild sensitivity to this parameter—making it an order of magnitude larger will lead to a more exploratory trajectory while making it an order of magnitude smaller will lead to less exploration. If λ is set to zero, no movement will occur at all. For further discussion of this point, see ***Miller et al.*** (***2016***). Finally, a sensitivity analysis is also provided in Fig. S6
Weight of the control term in the cost function of trajectory optimization loop (see Alg. S1)	R	10, 20	R is initially chosen to fit weakly electric fish behavior and kept the same for all the simulations except for moth, where R is set to 20 to account for the increased bandwidth of the sum-of-sine trajectory. Note that the control cost, when evaluated is equivalent to the total kinetic energy required to execute the candidate trajectory.
Number of dimensions used for Sobolev space norm in ergodic metric	d_S	15	d_S is initially chosen to be a sufficient number for representing all the behavioral data considered in this paper and kept the same for all the simulations
Initial control input	$\alpha(0)$	0	Zero control is applied at the beginning of every simulation
Initial belief	$p(\theta_0)$	unif$(0, 1)$	Initial belief is set and fixed to an uniform ("flat") prior distribution within the workspace (from 0 to 1) where the probability of the target located at every location is identical.

Table 2. Simulation Parameters Used For Each Figure

Figure	Category	SNR (dB)	Initial Position	Target Trajectory	Biological Condition
Figure 2D	Weak Signal	≤30	0.4	Sinusoid	N/A (simulation)
Figure 2E	Strong Signal	≥50	0.4	Sinusoid	N/A (simulation)
Figure 3A & 4B,D	Strong Signal	≥50	0.4	Prescribed by current study (0.1 Hz, 17 mm)	No jamming
Figure 3A & 4B,D	Weak Signal	≤30	0.4	Prescribed by current study (0.1 Hz, 17 mm)	≥ 10 mA jamming
Figure 3B & 4F,H	Strong Signal	≥50	0.2	Stationary	Intact control
Figure 3B & 4F,H	Weak Signal	≤30	0.6	Stationary	Single-side nares block & crossed airflow
Figure 3C & 4J,L	Strong Signal	≥50	0.475	Stationary	4 mm intact antenna
Figure 3C & 4J,L	Weak Signal	≤30	0.4	Stationary	1&2 mm bilaterally trimmed antenna
Figure 4N,P	Strong Signal	≥50	0.4	Prescribed by study ***Stockl et al.*** (***2017***)	3000 lux "high-light"
Figure 4N,P	Weak Signal	≤30	0.4	Prescribed by study ***Stockl et al.*** (***2017***)	15 lux "low-light"
Figure 5A-C & 6B	Weak Signal	≤30	0.4	(same as Figure 3A)	N/A (simulation)
Figure S1C	Weak Signal	10	0.4	(same as Figure 3A)	N/A (simulation)
Figure S1D	Strong Signal	60	0.4	(same as Figure 3A)	N/A (simulation)
Figure S2	Weak Signal	≤30	(same as Figure 3)	(same as Figure 3)	(same as Figure 3)
Figure S3	Strong and Weak Signal	10–55	0.4	(same as Figure 3A)	N/A (simulation)
Figure S4	Weak Signal	30	0.4	Stationary	N/A (simulation)
Figure S5	Weak Signal	30	0.4	Stationary	N/A (simulation)
Figure S6	Weak Signal	20	0.4	(same as Figure 3A)	N/A (simulation)
Figure S8	Strong Signal	≥50	0.8	Prescribed by study ***Khan et al.*** (***2012***)	Sham stitching
Figure S8	Weak Signal	≤30	0.3	Prescribed by study ***Khan et al.*** (***2012***)	Single-side nares stitching
Figure S9	Weak Signal	≤30	0.45	Stationary	Single-size nares block
Figure S9	Weak Signal	≤30	0.9	Prescribed by study ***Khan et al.*** (***2012***)	Single-side nares stitching

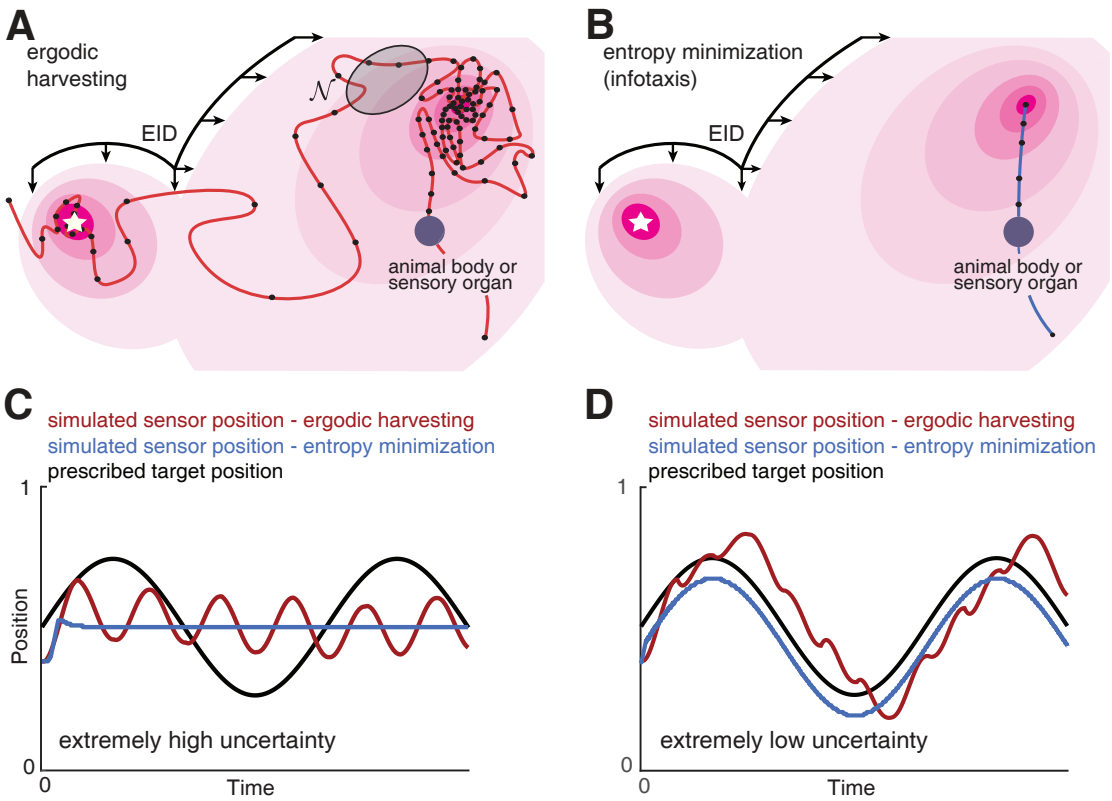

Figure S1. Illustration of differences between EIH and entropy minimization. (**A-B**) An illustration of an animal localizing a target of interest through moving its sensory organs or whole body with the proportional-betting (**A**) and information maximization (entropy minimization) (**B**) strategy. Black dots along the red trajectory indicate fixed time intervals, so longer distances between dots indicates higher speed. Uknown to the animal, the true target is the white star on the left, while on the right is a distractor (at center of red zone). The heat map represents the expected information density. Each step of EIH (Supplement Algorithm S1, **A**) is taken so that the resulting trajectory spends time in areas (*e.g.* $\mathcal{N}$) proportional to the expected information of that area. Trajectories generated by EIH, such as the one shown here, are ergodic with respect to the expected information because the percentage of time the trajectory spends in any subset $\mathcal{N}$ from $t = 0$ to $t = T$ is equal to the measure of expected information within $\mathcal{N}$; this condition must hold for all possible subsets. Information maximization (**B**), by contrast, locally maximizes the expected information (*i.e.* EID) at every step, leading it to go straight to the nearby distractor by following the gradient in this case. (**C-D**) Simulated animal tracking a target moving in 1D sine wave. Here we show simulated entropy minimization tracking trajectories alongside the corresponding simulated EIH trajectories under extreme SNR conditions (SNR of 10 dB and 55 dB, also shown in Fig. S3A). (**C**) With extremely high uncertainty (SNR $= 10$ dB), both algorithms were unable to accumulate enough information to successfully track the target. EIH results in persistent activation of movement even when the expected informativeness of sensing locations is flat due to a lack of information. The trajectory is prescribed to visit all locations of space with equal probability, covering the majority of the workspace (but not the entire workspace due to the energy constraint in trajectory optimization; see Supplement Algorithm S1). Entropy minimization results in cessation of motion under the same conditions. (**D**) Same as (C), but under extremely low uncertainty (strong signal) condition (SNR $= 60$ dB). Now entropy minimization hugs the one peak of the EID (there are two, as shown in Fig. 2C), while the EIH trajectory samples both peaks of the EID with excursions away from these peaks, similar to electric fish refuge tracking in the strong signal condition (Fig. 3A), and what can be inferred from a prior behavioral study in bat echolocation ***Yovel et al.*** (***2010***).

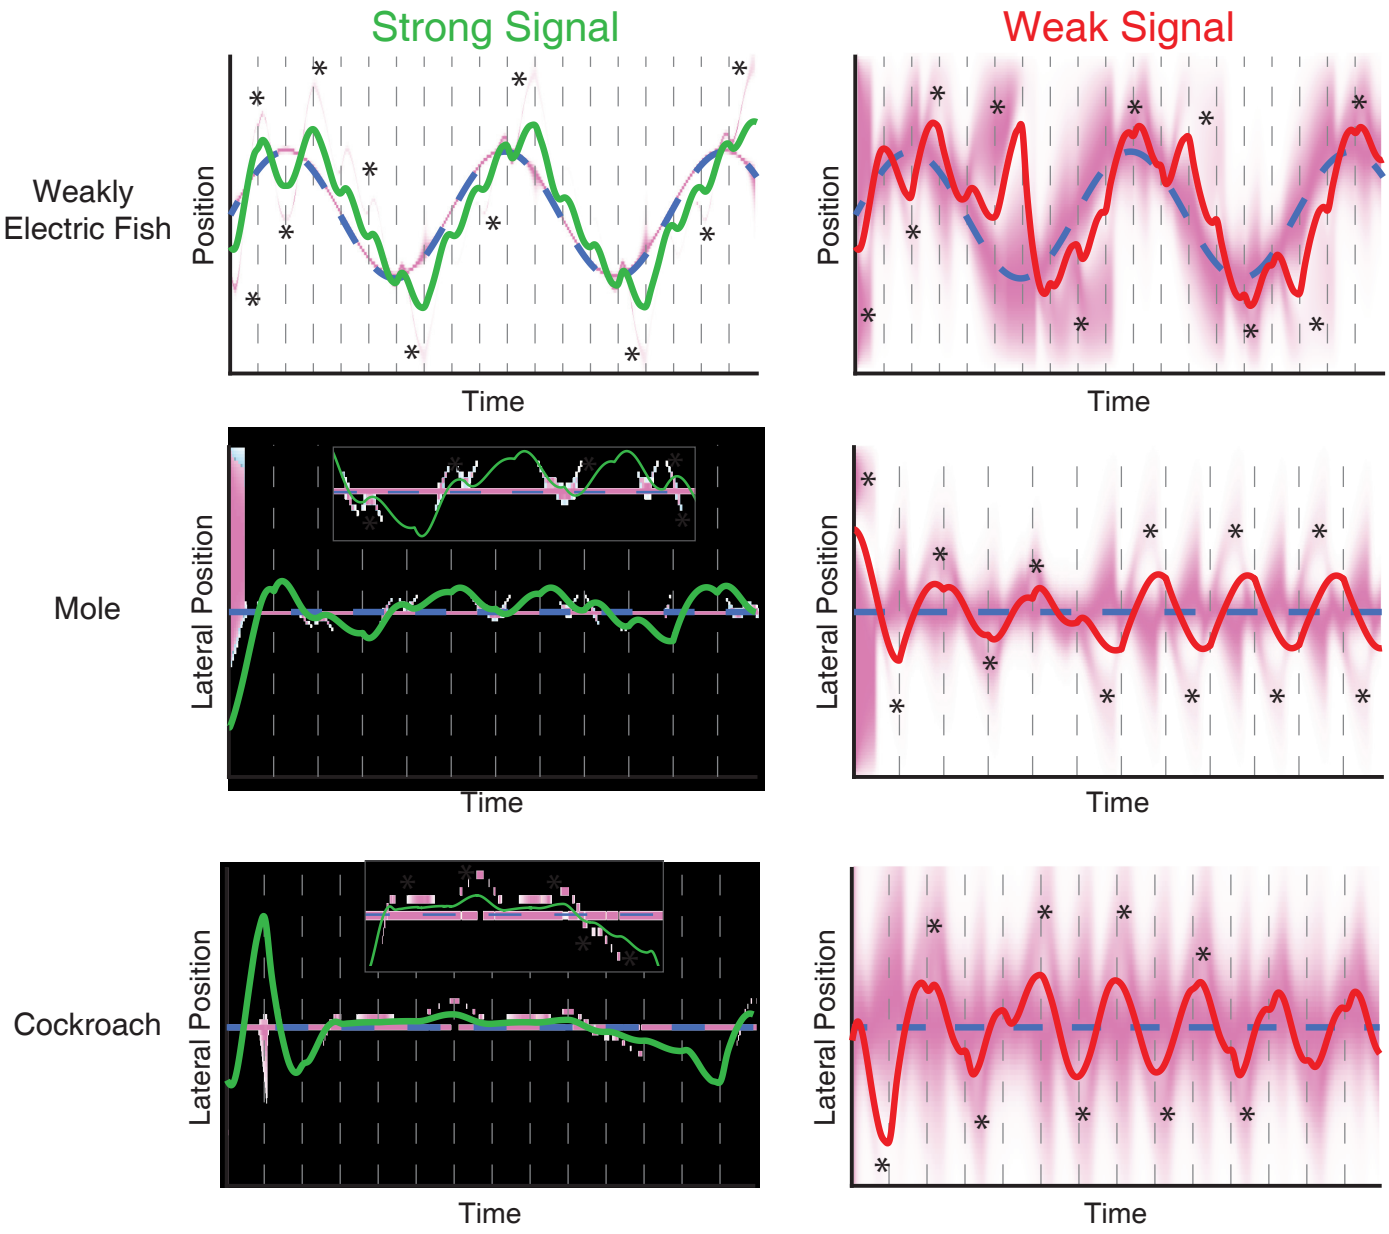

Figure S2. Evolution of belief over time for the trials shown in Fig. 3. EIH simulations of tracking behavior of weakly electric fish, mole, and cockroach. The trials shown are identical to those shown in Fig. 3. Simulated sensor position over time is the solid green line for the strong signal condition, and the solid red line in the weak signal condition. The target position is the dashed blue line. The repeated dashed gray lines mark each planning update event that segment the trajectory into individual planned trajectory segments of length T. The magenta heatmap shows the progression of the belief over time. Note the presence of fictive distractors, where the belief distribution becomes multi-peaked (asterisks). The inset in the mole and cockroach trials show a magnified segment to better visualize these distractors in the strong signal case.

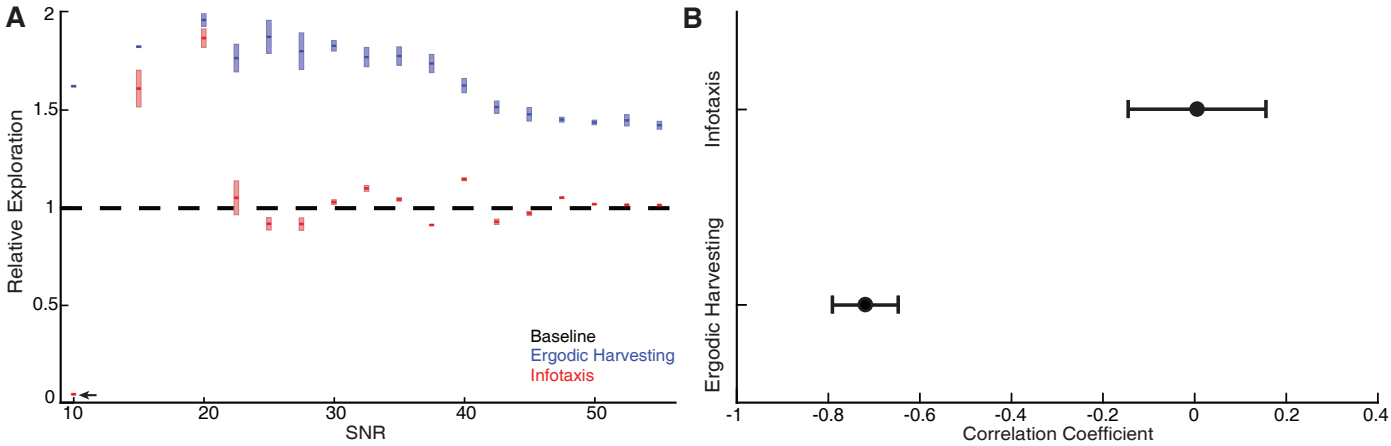

Figure S3. Systematic comparison between EIH and infotaxis in tracking a target moving in sinusoid. Simulations of sensor tracks a target moving in sinusoid under 17 different SNR conditions from 10 dB to 55 dB. For each SNR condition, 10 simulations with different pseudo-number seeds are performed to establish the confidence interval ($n = 170$ for EIH, $n = 170$ for infotaxis). (**A**) As the SNR decreases from 55 dB, EIH exhibits elevated relative exploration (see Methods) in the form of increased wiggle amplitude. This is consistent with the animal behaviors summarized in Figures 2-3. EIH's elevated exploration with decreasing SNR tapers at ≈20 dB as the Bayesian filter fails to converge the posterior due to significantly less informative measurements. Infotaxis does not show the trend of increased exploration as signal weakens between 55–20 dB, and prescribed cessation of movement at the lowest (10 dB) signal strength. Throughout most of the frequency spectrum, infotaxis attempts to hug one edge of the EID, leading to near unity relative exploration. (**B**) The correlation coefficient between relative exploration and signal strength computed from **A** for EIH and infotaxis (*n*=170 trials for EIH, and *n*=170 for infotaxis). The vertical line indicates the 95% confidence interval. Ergodic harvesting has a significantly more negative correlation, indicating exploration increases as signal strength declines. The same trend is absent in infotaxis.

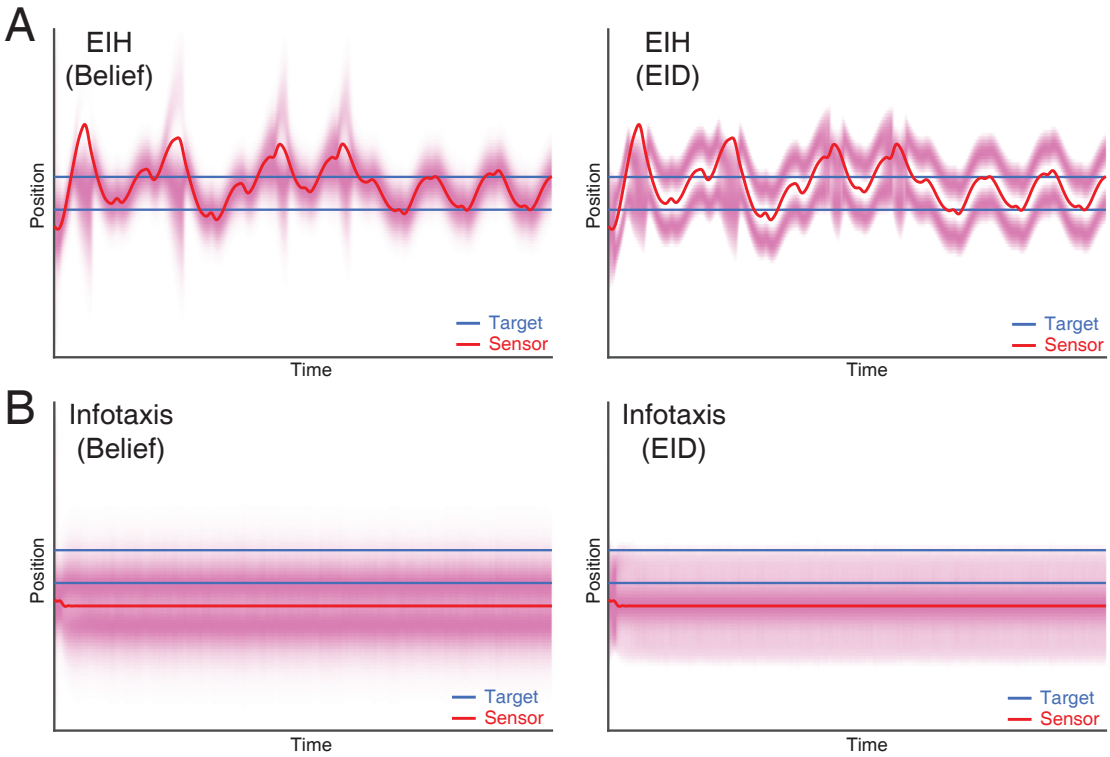

Figure S4. Dual target tracking simulation with EIH and infotaxis. In this simulation, two identical targets (with the same observation model) are present in the workspace, indicated by the two blue lines. To help visualize the outcome, both the belief (left panel) and EID (right panel) distribution over time are shown in magenta. In both cases, we use identical SNR levels corresponding to a weak signal condition (30 dB, see Table 2). (**A**) EIH tracks both targets with an oscillatory motion. In addition, belief peaks corresponding to the top and bottom targets are clearly visible in the left panel. (**B**) Infotaxis gets trapped at one of the information maximizing peaks, as seen in the right panel. This leads to cessation of movement and failure to detect the other target on the top (left panel).

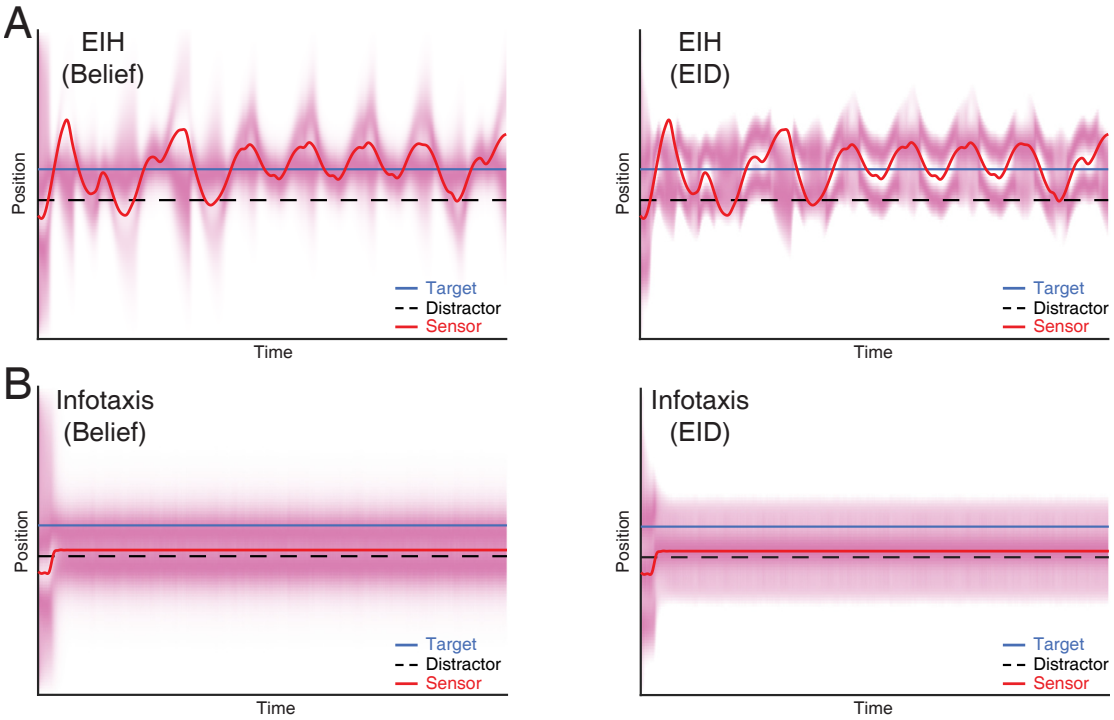

Figure S5. Single target tracking simulation with EIH and infotaxis in the presence of a simulated physical distractor. In this simulation, a single target and a physical distractor coexist in the workspace. The simulated physical distractor has a different observation model that leads to a different measurement profile when compared to the desired target. To help visualize the outcome, both the belief (left panel) and EID (right panel) distribution over time are shown in magenta. In both cases, we use identical SNR levels corresponding to a weak signal condition (30 dB, see Table 2). (**A**) EIH is able to find the desired target (the continuous peak along the desired target location in the left panel) and rejects the distractor (no continuous peak along the distractor location in the left panel). (**B**) Infotaxis gets trapped at one of the information maximizing peaks, as seen in the right panel. Although it does find the desired target, it fails to reject the distractor, as seen in the continuous peak near the distractor in the left panel.

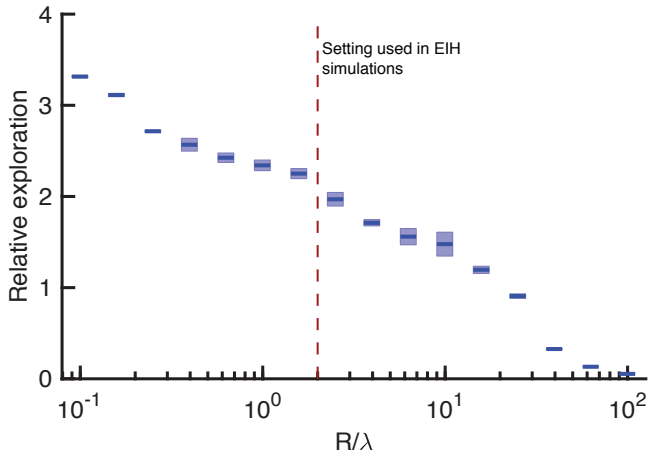

Figure S6. Sensitivity analysis on the ratio between control cost and ergodic cost in the objective function of trajectory optimization. Simulations are conducted in the same way as Fig. S3 but only for a fixed SNR under weak signal conditions (20 dB). The control cost term R is varied while fixing the ergodic cost term λ to be 5 for all the simulations. We simulated 16 different ratios of R/λ (shown along x axis on a log scale), each done with 10 different pseudo-random number seeds to establish confidence intervals. The relative exploration of the simulated trajectories are shown in the y axis on a linear scale with the sample mean marked by the solid color line and 95% confidence interval marked by the vertical color patch. The R/λ ratio used for all the EIH simulations included in the results is shown by the red dashed vertical line. Overall, as the R/λ drops, relative exploration drops monotonically because movement incurs higher cost. The trend also suggests that as the R/λ ratio increases, movement will cease. The value of $R/\lambda = 2$ that we used in our study results in a relative exploration value of around 2. The variation in relative exploration with an order of magnitude change in R/λ from 1 to 10 is 1.4 to 2.3. The same sensitivity analysis was also done for a strong signal condition but is not shown as the trend is similar.

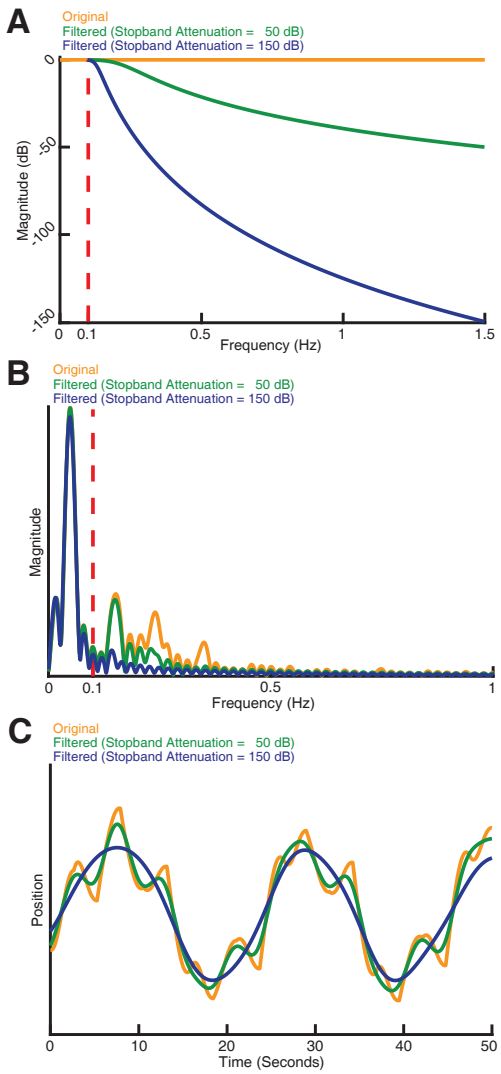

Figure S7. How sensor wiggles were attenuated for analyzing the impact of their removal. (**A**) Response of three types of wiggle attenuation filters. We used an IIR lowpass filter with a cutoff frequency of 0.1 Hz to avoid filtering the baseline tracking frequency band of the target (≈0.05 Hz in this case). In the two filtered examples shown, the attenuation at 1.5 Hz is set to 50 dB and 150 dB to achieve progressively stronger attenuation of sensor wiggling. (**B**) Frequency response of the original and filtered trajectory. The baseline tracking gain between 0 Hz to 0.1 Hz is well preserved while the wiggling frequency band between 0.1 Hz to 1.5 Hz is attenuated according to the attenuation gain. (**C**) Comparison of original and filtered trajectories. As the stopband attenuation increases, wiggle amplitude decreases.

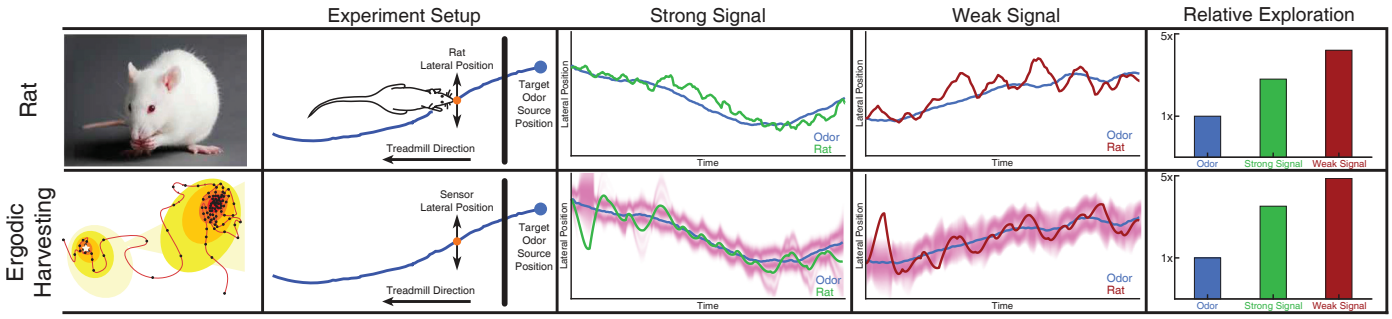

Figure S8. Rat odor tracking behavior and EIH simulation. In a prior study (***Khan et al., 2012***), Wistar rats (*Rattus norvegicus*, Berkenhout 1769) performed an odor tracking task by following a uniform odor trail on a moving treadmill with only olfactory information, under two different surgical conditions: with one side of the nose blocked (single-side nares stitching) and no blocking (sham stitching) as a control group. The experiment setup and tracked trajectory is shown in the top panel. Position data on the axis parallel to the treadmill were ignored, while positions on the perpendicular axis were recorded. The relative exploration was computed as with electric fish, using the cumulative 1D distance traveled by the rat divided by the cumulative distance of the transverse location of the odor trail. The olfactory system of the rat was modeled in the same fashion as for electric fish: as eventually providing an estimate of the location of the odor trail, modeled as a single 1D point-sensor of location with a Gaussian probability distribution controlled by the SNR of simulation (Methods). The blocked nares condition is here counted as the weak signal with lower SNR, and the sham stitching condition is counted as the strong signal. In the second row, we show EIH simulations of the weak and strong target trajectories, with the EID profile overlaid in magenta. It shows similar predicted sensor trajectories with a similar change in the relative exploration between these signal conditions as measured. These data are not included in the main study due to the low number of trials.

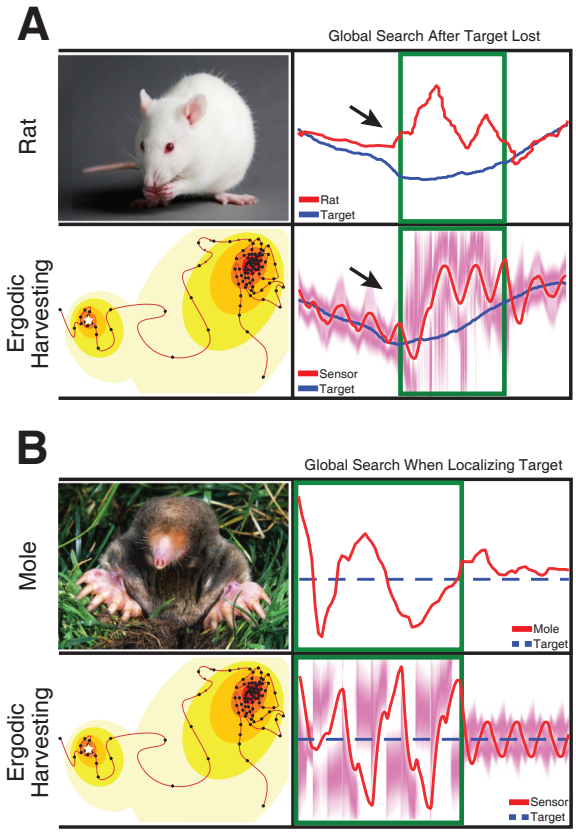

Figure S9. Measurements compared to prediction of EIH in two conditions where an animal needs to find the signal during tracking behavior. As a demonstration of how our model seamlessly transitions between exploitation and exploration, we examined an instance in the measured behavior of the live animals in which the rat appears to lose a signal it is following ***Khan et al.*** (***2012***), and an instance in which the mole starts off without finding the path to the source ***Catania*** (***2013***). The measured result shown here for both animals was wider casting motions transverse to the direction of movement. We examined whether these two behaviors (wider casting with loss of signal during tracking or when not finding the signal initially) would be predicted by EIH. (**A**) To simulate the condition of losing a signal trail in the middle of a successful tracking trial, we ran EIH as for the rat olfactory tracking case of Fig. S8, but with the trajectory shown in in panel A. In the "target lost" region (green square in figure), the sensor's input was replaced by random draws from the observation model. (**B**) To simulate the condition of starting without knowing where the trail is for an extended period of time, we ran EIH as for the mole olfactory tracking case of Fig. 3B, but with the trajectory shown in panel B. In the "target not yet acquired" region (green square in figure), the sensor's input was replaced by random draws from the observation model. In both instances, we observed wider casting motions transverse to the direction of travel similar to those measured (output from EIH below experimental data, overlaid with the EID profile in magenta). Similar to the measured behavior, our algorithm seamlessly transitions from exploitation to exploration when the expected information density becomes more diffuse.

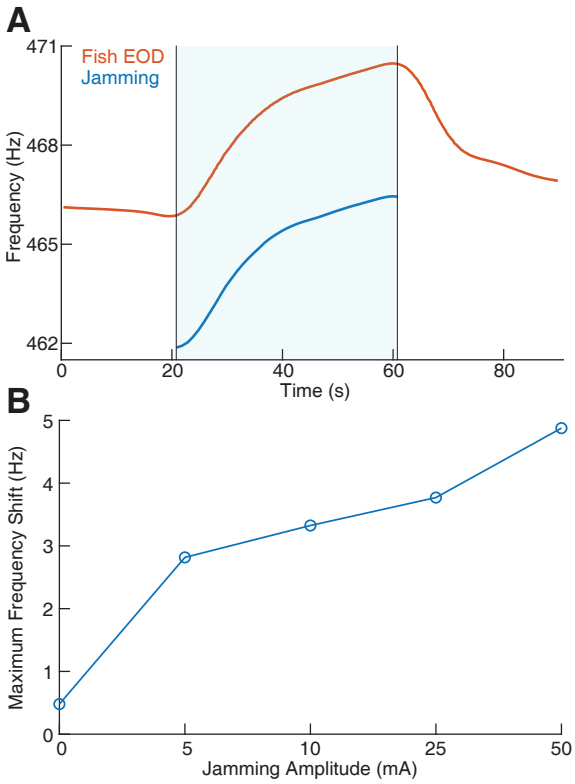

Figure S10. Effect of jamming and how it varies with jamming intensity. (**A**) The fish's electric organ discharge (EOD) frequency shifts up continuously as the jamming signal is being applied. The area shaded with light blue indicates when jamming is turned on and lasts 40 seconds. The jamming signal frequency (current amplitude = 25 mA) is set to be 5 Hz below the fish's EOD frequency. The fish's EOD frequency gradually declines back to its normal value after jamming is turned off after around 60 seconds. (**B**) Maximum frequency shift during the fixed 40 second jamming window as a function of jamming current amplitude. Maximum frequency shift is averaged for every jamming amplitude across a total of 26 trials ($n = 6$ for 0 mA, $n = 5$ for 5 mA, $n = 3$ for 10 mA, $n = 5$ for 25 mA, $n = 7$ for 50 mA). The amount of jamming applied has a clear positive correlation with the EOD frequency shift.

References

Catania KC. Stereo and serial sniffing guide navigation to an odour source in a mammal. Nat Commun. 2013;

Khan AG, Sarangi M, Bhalla US. Rats track odour trails accurately using a multi-layered strategy with near-optimal sampling. Nat Commun. 2012;

Lasdon L, Mitter S, Waren A. The conjugate gradient method for optimal control problems. IEEE Transactions on Automatic Control. 1967;

Miller LM, Silverman Y, MacIver MA, Murphey TD. Ergodic Exploration of Distributed Information. IEEE Transactions on Robotics. 2016 Feb;

Stockl AL, Kihlstrom K, Chandler S, Sponberg S. Comparative system identification of flower tracking performance in three hawkmoth species reveals adaptations for dim light vision. Philos Trans R Soc Lond, B, Biol Sci. 2017 Apr;

Thrun S, Burgard W, Fox D. Probabilistic robotics. MIT press

Yovel Y, Falk B, Moss CF, Ulanovsky N. Optimal localization by pointing off axis. Science. 2010;