Research

Author for correspondence:
L. Mahadevan
e-mail: lm@seas.harvard.edu

THE ROYAL SOCIETY
PUBLISHING

Optimal switching between geocentric and egocentric strategies in navigation

O. Peleg[1] and L. Mahadevan[1,2,3]

[1]Paulson School of Engineering and Applied Sciences, [2]Department of Physics, and [3]Department of Organismic and Evolutionary Biology, Kavli Institute for NanoBio Science and Technology, Wyss Institute for Biologically Inspired Engineering, Harvard University, Cambridge, MA 02138, USA

OP, 0000-0001-9481-7967

Animals use a combination of egocentric navigation driven by the internal integration of environmental cues, interspersed with geocentric course correction and reorientation. These processes are accompanied by uncertainty in sensory acquisition of information, planning and execution. Inspired by observations of dung beetle navigational strategies that show switching between geocentric and egocentric strategies, we consider the question of optimal reorientation rates for the navigation of an agent moving along a preferred direction in the presence of multiple sources of noise. We address this using a model that takes the form of a correlated random walk at short time scales that is punctuated by reorientation events leading to a biased random walks at long time scales. This allows us to identify optimal alternation schemes and characterize their robustness in the context of noisy sensory acquisition as well as performance errors linked with variations in environmental conditions and agent–environment interactions.

1. Introduction

Navigation in complex uncertain environments requires information about environmental cues (landmarks) along with the ability to memorize and execute intended plans based on these cues. It is thus often accompanied by several cognitive demanding activities, such as multi-sensory acquisition and integration, locomotion planning and motor control. As organisms have finite cognitive and computational resources, they must therefore multitask and develop optimal schemes to dynamically allocate resources to different tasks [1–6]. Similar demands are also placed on their artificial analogues, autonomous vehicles such as robots, self-driving cars and even deep-space craft [7,8]. Typical strategies for navigation involve a combination of egocentric and geocentric schemes. In the first, the organism uses information that it has

acquired somehow to move along a path without any realtime feedback. In the second, the organism constantly probes its location relative to environmental cues and adjusts its strategy in real time. The first requires perfect memory and can be efficient, but can succeed only in a constant environment; the second requires continuous course corrections and is often slow, but can cope with fluctuating environments. Organisms switch between these strategies to navigate accurately in a variable environment, and understanding the switching strategies of organisms in their natural environments is a basic problem in neuroethology.

Dung beetles provide a well-studied example of this switching behaviour during navigation [9]. Foraging beetles look for nutrient-rich dung, and then attempt to roll a dung ball along a straight path radially away from the pile [1] with the aim of providing food for their brood. Beetles acquire navigation information using a variety of long-range cues before initiating a roll, and then push the dung ball while walking backwards with their hind legs in contact with the dung. However, they stop intermittently and get atop the ball and walk on it to reorient themselves before continuing to roll the dung ball. The reorientation behaviour of the beetle is visually mediated and triggered both by active and passive deviations off-course and by visual cues [10] that include the position of the Sun, the Moon and associated light polarization patterns. While the beetle continues to obtain navigational information during a roll, the lack of a constant deviation angle triggering a reorientation behaviour suggests that this event has uncertainty [10]. Not having to rely on terrestrial landmarks allows the beetle to travel arbitrarily far away from its starting point without large directional errors [11], thus making the nature of the navigational task fundamentally different from that of homing insects. Furthermore, the switching behaviour between runs and reorientations is a function of the environment; the more uncertain the environment as characterized by its physical roughness, the more frequent the reorientations [10]. Stopping to reorient after each step makes for slow progress, but guarantees a straighter route. Alternatively, not stopping at all avoids the loss of time associated with reorientation, but leads to trajectories that deviate significantly from the intended bearing. This leads to a natural question of the optimal rate of switching between the egocentric strategy and the geocentric one. In this paper, we introduce a model for the movement of the beetle, or a navigational agent, that is associated with paths that are a characteristic of correlated diffusion on short time scales, and biased diffusion on long time scales [12]. This allows us to pose and solve an optimization problem for the most efficient switching strategy and characterize its robustness in the presence of noisy sensory acquisition and in rough environments, with relevance to questions that go beyond the original motivation for the problem.

2. Mathematical model

Inspired by the beetle, we consider an agent that performs a random walk interspersed by reorientation events, in which its heading direction is reset. We expect that owing to finite cognitive and attention capacity, information gathered while rolling the ball leads to a bigger acquisition error relative to that when it is reorienting, i.e. the detected orientation is $\theta_i + \epsilon$; we assume that this dynamic acquisition error, ϵ, is drawn from a uniform distribution of headings $[0, \epsilon_d]$, where $\epsilon_d \in [0, \pi]$. Between reorientation events, we assume that the random walk is correlated orientationally, with a mean corresponding to the current orientation and a uniform distribution of angular errors in the range $[0, \theta^*]$, where $\theta^* \in [0, \pi]$. The agent can acquire (visual) sensory information in two modes: when it is rolling the ball, or when it reorients. Memory degradation, execution errors and environmental uncertainties translate to larger accumulated errors, i.e. large θ^*, so that the agent must stop to reorient more often, consistent with observations of beetles [10]. This suggests that reorientation is triggered by a threshold deviation from the preferred heading, which we denote by θ_a, where $\theta_a \in [0, \pi]$. However, reorientation events do not occur at the same angular deviation from the original bearing [10], and therefore requires the addition of further triggering mechanisms. Indeed, beetles detect gradual deviation better than abrupt ones [10]. This suggests that beetles do not pay attention to navigational signals at all times owing to finite cognitive resources; when paying attention to navigational cues, motor control suffers and vice versa. To quantify this, we define the attention span, τ, where $\tau \in [0, 1]$ is the fraction of time during movement when the agent pays more attention to navigational cues, with a distribution of turning angles drawn from a uniform distribution $[0, A\theta^*]$ (with $A > 1$), while over the fraction of time $1 - \tau$, the distribution of angles is in the range $[0, \theta^*]$. If the agent is inattentive to navigation cues during a step, the standard deviation of the distribution of angles is θ^*, else it increases by a factor $A > 1$ and becomes $A\theta^*$ over the attentive steps $(1 - \tau)N$. In other words, when paying attention to the navigational cues, the agent pays less attention to its motor control and the resulting trajectory is characterized by a random walk with a
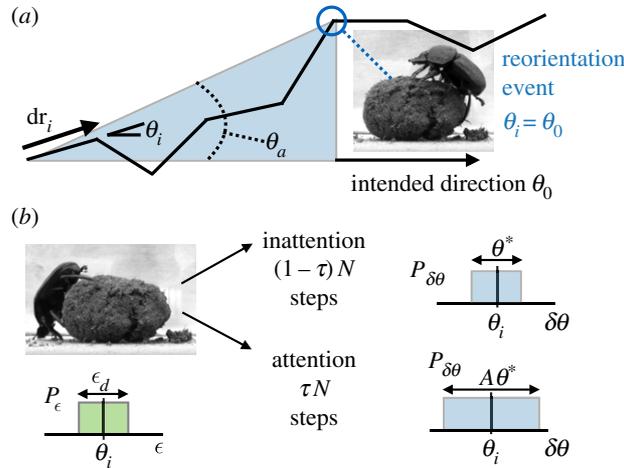
Figure 1. (a) Schematic of the navigation model for the dung beetle. The beetle (agent) walks along a correlated random walk, with errors introduced at every step. When the accumulated error exceeds a threshold θ_a, the (agent) resets its orientation and continues. (b) Strategies and associated errors. If the agent is inattentive to navigation cues during a step, the standard deviation of the distribution of angles is θ^*, else it increases by a factor $A > 1$ and becomes $A\theta^*$ over the attentive steps $(1 - \tau)N$. Similarly, there is an error associated with sensory acquisition that is assumed to be uniformly distributed with a standard deviation ϵ_d. Dung beetles pictures acquired by the Vision Group in Lund University [13].

wider turning angle distribution. These sources of stochasticity in sensing, attention and movement as the agent interacts with the environment are depicted in figure 1.

The agent is assumed to start at the origin and walk in a straight line along a given radial direction θ_0. At each time step i the agent moves in a random direction θ_i (figure 1a), relative to its current heading, i.e. $\theta_i = \theta_{i-1} + \delta\theta$, where for simplicity $\delta\theta$ is drawn from a uniform distribution $[0, \theta^*]$ [12] (choosing from a Gaussian distribution does not change any of results qualitatively). The agent stops at certain intervals to reorient, i.e. it resets $\theta_i = \theta_0$, which we assume takes time, given that the agent has to take a bearing (corresponding to the beetle reorienting on the dung ball). Then, after N steps that include N_r reorientations, the end–to–end vector of the agent is $\mathbf{R} = \sum_{i=1}^{N} \mathbf{dr}_i$, where $\mathbf{dr}_i = (\sin\theta_i, \cos\theta_i)$.

To characterize the competition between accuracy and speed associated with the task of moving as far from the origin as quickly as possible, we define a simple cost function that penalizes the deviation from a straight line and also penalizes the time spent for reorientation:

$$f = \frac{N - |\mathbf{R}|}{N} + \frac{N_r}{N}. \tag{2.1}$$

Frequent reorientations result in $|\mathbf{R}| \to N$ and leads to the first term being small, but the second being large. By contrast, few reorientations will make the second term small, while the first will be large. Continuity suggests that the cost will be minimum when the number of steps before a reorientation event $n = N/N_r$, or equivalently, the frequency of reorientation, is an optimum. Before proceeding to find this, we point out that this problem has some similarities to a recent class of problems named 'search with reset' [14–17], but differs qualitatively in that here we consider diffusion in orientational rather than rectilinear space, while coupling reorientation to translational motion, making an analytic approach difficult except in the simplest of situations.

3. Analysis of model and results

3.1. Regular reorientation

We start with the simplest variant of the question of optimal switching, assuming that the heading direction is reset to zero every n steps, with $N = nN_r$, complete inattention to navigational cues between reorientation events, i.e. $\tau = 0$, so that there is no error amplification, i.e. $A = 1$, and finally there is no acquisition error, i.e. $\epsilon_d = 0$. Later, we will include the additional stochasticity associated with randomness in the choice of reorientation intervals, and errors in acquisition, variable attention span and accompanying error amplification.
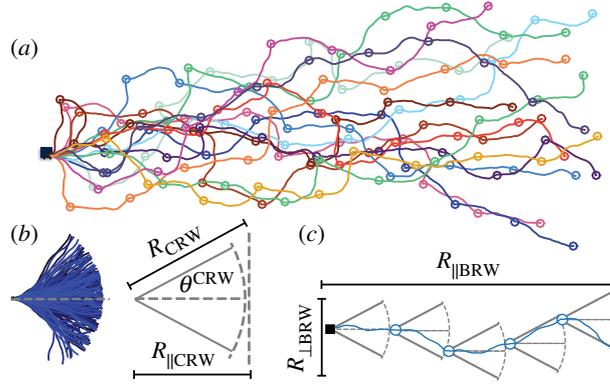
Figure 2. (*a*) Examples of trajectories obtained via random walk simulations with fixed reorientation interval. Different realizations are shown in different colours and reorientation points are shown as empty circles. The origin is marked with a black square. $\theta^* = \pi/12$, $N = 100$ and $n = 20$. (*b*) CRWs of $n = 20$ steps and $\theta^* = \pi/12$. Each realization is shown in a different shade of blue. The area accessible to the agent between reorientation events is shown as a circular sector. (*c*) The entire walk is described as a BRW with step size and turning angles defined by the CRW statistics.

With these assumptions, the agent's path is a correlated random walk (CRW) in between reorientation events, and a biased random walk at the long time scale (BRW) [12], as illustrated in figure 2*a*,*b*. The path always starts with a step along θ_0, followed by an $n-1$ steps CRW, resulting in a mean square end-to-end distance $\langle R_{CRW}^2 \rangle$ and mean deviation angle θ^{CRW}. These two quantities define a sector accessible to the agent between reorientation events with central angle $2\theta^{CRW}$ and radius $\sqrt{\langle R_{CRW}^2 \rangle}$ (figure 2*c*). Then, following [11]:

$$\langle R_{CRW}^2 \rangle = n - 1 + \frac{2\beta}{1-\beta}\left(n - 1 - \frac{1-\beta^{n-1}}{1-\beta}\right) \tag{3.1}$$

$$\langle R_{\|CRW} \rangle = \frac{\beta - \beta^n}{1-\beta} \tag{3.2}$$

and

$$\theta_{CRW} = \cos^{-1}\left(\frac{\langle R_{\|CRW} \rangle}{\sqrt{\langle R_{CRW}^2 \rangle}}\right), \tag{3.3}$$

where n is the number of steps before a reorientation, and $\beta = \langle \cos\theta_i \rangle$, $\langle R_{\|CRW} \rangle$ is the mean distance travelled along the directional bias and θ_{CRW} is the mean deviation from the intended direction.

On longer time scale, the entire random walk may be treated as a BRW with individual steps given by the CRWs (figure 2*c*). The end-to-end distance of a BRW is defined as $\langle R_{BRW}^2 \rangle = \langle R_{BRW\|}^2 \rangle + \langle R_{BRW\perp}^2 \rangle$, where $\langle R_{BRW\|}^2 \rangle$ and $\langle R_{BRW\perp}^2 \rangle$ are the mean square displacements parallel and perpendicular to the directional bias, respectively. Following [11]:

$$\langle R_{BRW\perp}^2 \rangle = N_r(1 - \gamma_{BRW})\langle dr_{BRW}^2 \rangle \tag{3.4}$$

and

$$\langle R_{BRW\|}^2 \rangle = N_r^2 \gamma_{BRW}\langle dr_{BRW}^2 \rangle + N_r^2, \tag{3.5}$$

where $\gamma_{BRW} = \langle \cos^2\theta_{CRW} \rangle$ and the mean square step size $\langle dr_{BRW}^2 \rangle = \langle R_{CRW}^2 \rangle$. Here, we have added the N_r term to $\langle R_{BRW\|}^2 \rangle$ to account for the fact that the first step after reorientation is parallel to the directional bias. Therefore, the mean square displacement of the entire walk is:

$$\langle R_{BRW}^2 \rangle = N_r(1 - \gamma_{BRW})\langle R_{CRW}^2 \rangle + N_r^2 \gamma_{BRW}\langle R_{CRW}^2 \rangle + N_r^2. \tag{3.6}$$
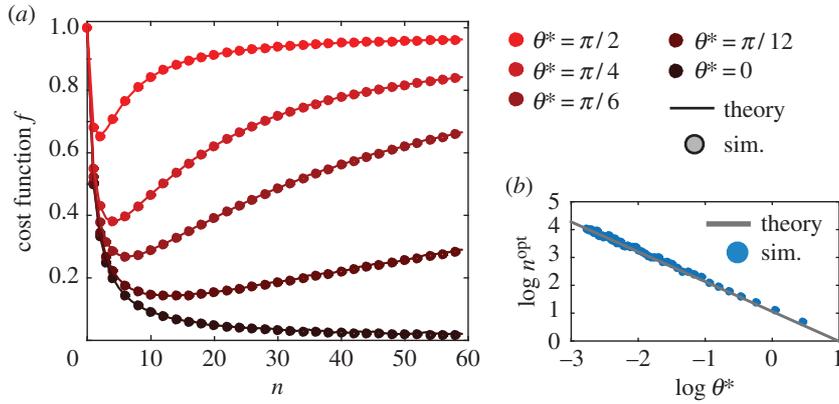
Figure 3. (a) Cost function versus n, for simulations with 5000 different realizations, $N = 1000$. The curves are coloured according to their θ^* value (see legend). Theoretical prediction shown as solid lines and numerical values obtained via simulations shown as open circles. (b) Log–log plot of $n^{opt}(\theta^*)$ versus θ^* derived from both the theoretical prediction (red solid line), and simulations (open circles).

Using (3.1) and (3.3) in the above relationship allows us to calculate $|\mathbf{R}| = \langle R^2_{BRW} \rangle^{1/2}$, so that the cost given in (2.1) becomes:

$$f(N, n, \beta(\theta^*)) = \frac{N - \sqrt{\langle R^2_{BRW} \rangle}}{N} + \frac{1}{n}$$

$$= 1 + \frac{1}{n} - \frac{\sqrt{N}}{N(\beta - 1)n} \sqrt{N(\beta^n - 2\beta + 1)^2 - (\beta^2 - 1)n^2 - n(\beta^n - 1)(\beta^n - 2\beta - 1)} \quad (3.7)$$

In figure 3, we show the dependence of the cost given by (3.7) on the frequency of reorientation n. When the width of the turning angle distribution $\theta^* = 0$, the agent travels along a perfectly straight line and the cost function f decreases monotonically with n (black solid curve). For a uniform turning angle distribution, $\beta = \sin\theta^*/\theta^*$, so that when $\theta^* \neq 0$, f is optimal for a particular $n(\theta^*$ (coloured solid curves). We confirm these analytical results using simulations, as shown in figure 3a.

To study how the optimal reorientation frequency, n^{opt} depends on θ^*, we consider the cost function in the asymptotic limit $f(N \to \infty, n, \beta(\theta^*))$:

$$\lim_{N \to \infty} f = \frac{(\beta - 1)(1 + n) + 1 - \beta^n}{(\beta - 1)n}. \quad (3.8)$$

Minimizing this with respect to n, yields

$$n^{opt} = \frac{1 + W(-\beta/e)}{-\log\beta}, \quad (3.9)$$

where W is the Lambert W function (the solution of $z = W \exp W$). For the case when the turning angle distribution is uniform:

$$n^{opt}(\theta^*) = \frac{1 + W(-\sin\theta^*/\theta^*e)}{-\log(\sin\theta^*/\theta^*)} \sim \frac{\pi}{\theta^*}, \quad (3.10)$$

(see electronic supplementary material for a derivation). In figure 3b, we plot $n^{opt}(\theta^*)$ versus θ^* derived from our simulations, and see that they agree well with our simple theoretical prediction. Furthermore, figure 3a shows that as θ^* increases it becomes more crucial to be close to the optimal reorientation interval, as small deviations in n result in a sharp increase of the cost.

3.2. Dynamic reorientation

Having understood this simple scenario, we now go back to address the complexities associated with noise in acquisition, planning and execution, in addition to the noise in the turning angle distribution θ^*. To address this, we use numerical simulations with the acquisition error $\epsilon_d \neq 0$, the error amplification associated with attention $A > 1$, and ask how the agent must optimize the attention span $\tau \neq 1$ and the threshold activation angle for reorientation θ_a to minimize the cost function defined in equation (2.1). We use the covariance matrix adaptation algorithm (CMA–ES) [18] to determine the optimal strategy, where given a set $(\theta^*, A, \epsilon_d)$, we determine the optimal set (θ_a, τ). The CMA-ES is a stochastic derivative-free
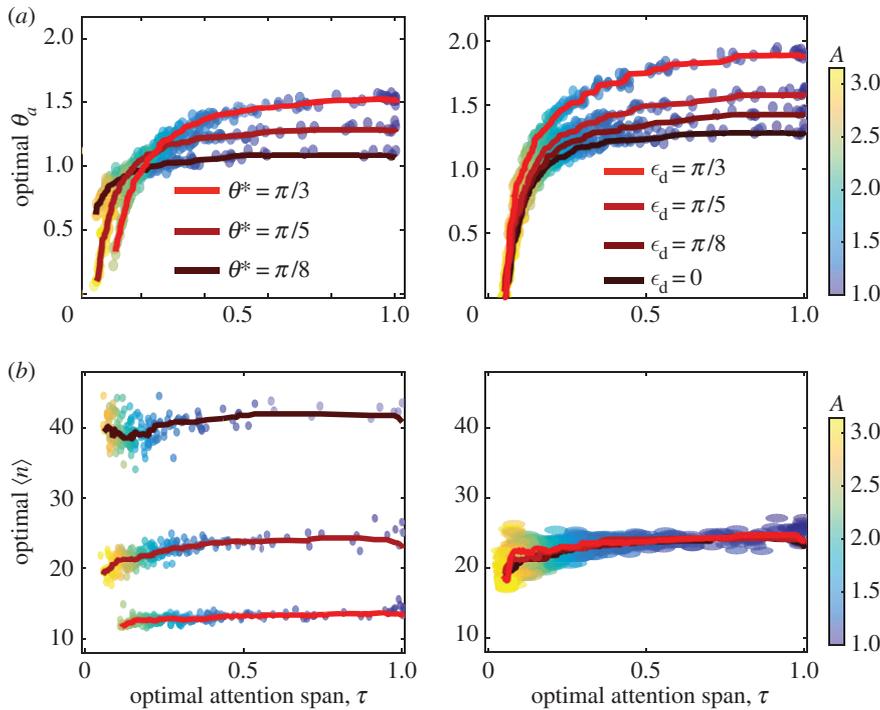
Figure 4. Optimal strategies in stochastic navigation. (*a*) Optimal activation angle θ_a and attention span τ. (*b*) The optimal mean reorientation interval, $\langle n \rangle$, versus τ. Panels (*a,b*) summarize the effect of θ^* and ϵ_d, respectively on θ_a, τ. Each point in the scatter plots corresponds to a set of parameters controlled by the agent or the environment, as summarized in figure 1. Colour scale corresponds to the attention error amplification, A. Solid lines created using a Gaussian filter over the scatter points as shown in legend. $N = 100$; $\lambda_{CMA} = 10$ (population size per generation).

optimization method for nonlinear or non-convex continuous optimization problems. By incrementally increasing the probability of previously successful candidate solutions, we iteratively perform the following three steps: (i) sample p new sets of (θ_a, τ) following the distribution with the updated mean and covariance of the cost function f (with step size built in), (ii) evaluate f and re-order the sampled solutions based on their fitness, and (iii) update the internal state variables (θ_a, τ), including the mean, the isotropic and anisotropic evolution path, the covariance matrix, and the step size, based on the q best out of p solutions, until we have converged to the minimal cost that will yield $\tau = \tau(\theta^*, A, \epsilon_d)$, $\theta_a = \theta_a(\theta^*, A, \epsilon_d)$.

The results of our optimization are summarized in figure 4. We see that both the optimal attention span τ and optimal activation angle θ_a decrease monotonically with increase in the attention error amplification, A, consistent with our intuition. However, the mean reorientation interval, $\langle n \rangle$, (extracted retrospectively from the simulations), is relatively independent of A (figure 4*b*). When the variance of the turning angles θ^* increases, we see two regimes; for large A, τ is large and θ_a increases with θ^*. For small values of A, there is no clear trend. Finally, θ_a and τ increase monotonically with ϵ_d.

4. Discussion

Inspired by the behavioural strategy of navigational movement of the dung beetle, we have posed and solved an optimization problem of geocentric navigation interspersed by egocentric cue integration. In the simplest setting, we find that the optimal reorientation interval is inversely proportional to the environmental noise and is invariant to sensory acquisition noise. In more complex settings, our study highlights the variations in the optimal navigation strategy that balances accuracy, speed and effort. Our study might well be generalizable more broadly to animal navigation, and perhaps even artificial vehicular navigational situations, all of which use egocentric and geocentric cues with varying attention to create optimal strategies that balance accuracy and efficiency in the presence of noise, and additional constraints such as finite cognitive bandwidth. A natural next step would be to test the theory with

animals, e.g. dung beetles (using established experimental methods [19]), as well as autonomous robotic agents.

References

1. Warrant E, Dacke M. 2010 Visual orientation and navigation in nocturnal arthropods. *Brain Behav. Evol.* **75**, 156–173. (doi:10.1159/000314277)

2. Wehner R, Michel B, Antonsen P. 1996 Visual navigation in insects: coupling of egocentric and geocentric information. *J. Exp. Biol.* **199**, 129–140. (doi:10.5167/uzh-634)

3. Wiltschko W, Wiltschko R. 2005 Magnetic orientation and magnetoreception in birds and other animals. *J. Comp. Physiol. A* **191**, 675–693. (doi:10.1007/s00359-005-0627-7)

4. Papastamatiou YP, Cartamil DP, Lowe CG, Meyer CG, Wetherbee BM, Holland KN. 2011 Scales of orientation, directed walks and movement path structure in sharks. *J. Anim. Ecol.* **80**, 864–874. (doi:10.1111/j.1365-2656.2011.01815.x)

5. Balkovsky E, Shraiman BI. 2002 Olfactory search at high Reynolds number. *Proc. Natl Acad. Sci. USA* **99**, 12 589–12 593. (doi:10.1073/pnas.192393499)

6. Vergassola M, Villermaux E, Shraiman BI. 2007 'Infotaxis' as a strategy for searching without gradients. *Nature* **445**, 406–409. (doi:10.1038/nature05464)

7. Wood LJ. 2008 The evolution of deep space navigation: 1989–1999. In *Proc. of the 31st annual AAS Rocky Mountain Guidance and Control Conference*, Breckenridge, CO.

8. Quadrelli MB, Wood LJ, Riedel JE, McHenry MC, Aung M, Cangahuala LA, Volpe RA, Beauchamp PM, Cutts JA. 2015 Guidance, navigation, and control technology assessment for future planetary science missions. *J. Guid. Control Dynam.* **38**, 1165–1186. (doi:10.2514/1.G000525)

9. Byrne M, Dacke M. 2011 The visual ecology of dung beetles. In *Ecology and evolution of dung beetles* (eds LW Simmons, TJ Ridsdill-Smith), pp. 177–199. Chichester, UK: John Wiley & Sons, Ltd.

10. Baird E, Byrne MJ, Smolka J, Warrant EJ, Dacke M. 2012 The dung beetle dance: an orientation behaviour? *PLoS ONE* **7**, e30211. (doi:10.1371/journal.pone.0030211)

11. Cheung A, Zhang S, Stricker C, Srinivasan MV. 2007 Animal navigation: the difficulty of moving in a straight line. *Biol. Cybern.* **97**, 47–61. (doi:10.1007/s00422-007-0158-0)

12. Codling EA, Plank MJ, Benhamou S. 2008 Random walk models in biology. *J. R. Soc. Interface* **5**, 813–834. (doi:10.1098/rsif.2008.0014)

13. Lund Vision Group. 2012 Dung beetle dance. See https://youtu.be/w1XL711eIDA.

14. Bénichou O, Coppey M, Moreau M, Voituriez R. 2007 Intermittent search strategies: when losing time becomes efficient. *Europhys. Lett.* **75**, 349–354. (doi:10.1209/epl/i2006-10100-3)

15. Moreau M, Bénichou O, Loverdo C, Voituriez R. 2009 Chance and strategy in search processes. *J. Stat. Mech.* **2009**, P12006. (doi:10.1088/1742-5468/2009/12/P12006)

16. Evans MR, Majumdar SN. 2011 Diffusion with stochastic resetting. *Phys. Rev. Lett.* **106**, 160601. (doi:10.1103/PhysRevLett.106.160601)

17. Evans MR, Majumdar SN. 2011 Diffusion with optimal resetting. *J. Phys. A: Math Theor.* **44**, 435001. (doi:10.1088/1751-8113/44/43/435001)

18. Hansen N. 2006 An analysis of mutative sigma-self-adaptation on linear fitness functions. *Evol. Comp.* **14**, 255–275. (doi:10.1162/evco.2006.14.3.255)

19. Martínez-García R, Calabrese JM, Mueller T, Olson KA, López C. 2013 Optimizing the search for resources by sharing information: Mongolian gazelles as a case study. *Phys. Rev. Lett.* **110**, 248106. (doi:10.1103/PhysRevLett.110.248106)