

Learning Foraging Thresholds for Lizards: An Analysis of a Simple Learning Algorithm¹

Leslie Ann Goldberg, University of Warwick²

William E. Hart, Sandia National Laboratories³

David Bruce Wilson, Microsoft⁴

ABSTRACT This paper gives a proof of convergence for a learning algorithm that describes how anoles (lizards found in the Caribbean) learn a *foraging threshold* distance. An anole will pursue a prey if and only if it is within this threshold of the anole's perch. The learning algorithm was proposed by Roughgarden and his colleagues. They experimentally determined that this algorithm quickly converges to the foraging threshold that is predicted by optimal foraging theory. We provide analytic confirmation that the optimal foraging behavior as predicted by Roughgarden's model can be attained by a lizard that follows this simple and zoologically plausible rule of thumb.

1. Introduction

Anolis lizards are lizards found in the Caribbean that use a “sit-and-wait” predation strategy. Anoles usually perch at a spot in their territory scanning the ground for prey. If an anole sees a prey that it considers worth pursuing, it leaves its perch and returns after the prey has been caught.

¹ Part of this work was performed at Sandia National Laboratories and was supported by the U.S. Department of Energy under contract DE-AC04-94AL85000.

² Department of Computer Science, University of Warwick, Coventry CV4 7AL, United Kingdom. leslie@dcs.warwick.ac.uk. This work was partially supported by ESPRIT LTR Project no. 20244 — ALCOM-IT and Project no. 21726 — RAND-II and by EPSRC grant GR/L60982.

³ Sandia National Laboratories, PO Box 5800, Albuquerque, NM 87185-1110 USA. wehart@cs.sandia.gov.

⁴ Microsoft Corporation 31/3348, One Microsoft Way, Redmond, WA 98052. dbwilson@alum.mit.edu.

Optimal foraging [4,6] theory offers one approach to modeling the foraging behavior of anoles. Optimal foraging theory begins with the choice of an optimization criterion, which characterizes what it is that a lizard's foraging behavior optimizes. Roughgarden [11] suggests that a lizard may minimize the average time used to capture a prey, which is reasonable in circumstances of high predation. Roughgarden notes that minimizing the average time to capture a prey leads to a tradeoff for the lizard:

...if a lizard chases a very distant item, it is away from its perch and cannot see (or react) to prey that may appear while it is gone. Alternatively, a lizard may ignore a very distant item, and yet nothing may actually appear during the time it would have chased down that item. So, where should the lizard draw the line? Clearly it should *not* chase extremely distant prey, for it would be away from its perch all day, and it *should* chase extremely close prey.

Roughgarden computes the optimal foraging behavior of a lizard in a simple model in which the lizard pursues all prey that are within a foraging threshold. This analysis predicts an optimal foraging threshold, X^* ; if the lizard only pursues prey that land within distance X^* of its perch, then the average time that the lizard uses to capture its prey is minimized.

According to this analysis, X^* is the foraging threshold that the lizard *should* use. A natural question is *how* a lizard determines this foraging threshold [11, 3, 13]. Roughgarden argues that it is unlikely that a particular foraging threshold (i.e., a particular distance) is hard-wired into the brain of the lizard, because the rate a at which prey appear and the speed v of the lizard, which both affect X^* , depend upon where the lizard lives. Any particular foraging threshold might be close to optimal if the lizard lives in some places, and very far from optimal if it lives in other places. On the other hand, it is unlikely lizards have the cognitive ability to compute X^* from a and v since this involves calculating cube roots [11].

Roughgarden [11] argues that a more realistic model of foraging is that a lizard learns an appropriate foraging threshold based on its past experiences pursuing prey. He proposes a learning algorithm that continually refines the lizard's current foraging threshold. This algorithm assumes that the lizard can store two quantities: the time spent so far, and the total number of prey that have been captured.[†] Shafir and

[†] Shafir and Roughgarden [11, 13] actually divide the total time into two parts: the total time spent pursuing and the total time spent waiting for prey. To execute their learning algorithm, however, the lizard really only needs to keep track of the total time.

Roughgarden [13] note that experimental evidence suggests that many animals can conceptualize this type of information.

Experimental results with this learning algorithm are reported by Shafir and Roughgarden [11, 13]. These results confirm that the optimal foraging threshold is found very quickly by the learning algorithm. Menczer, Hart and Littman [8] prove that if the learning algorithm converges, then it converges to X^* . While this provides confirmation that the learning algorithm is reasonable, it does not constitute a proof of convergence.

Our analysis of Roughgarden's learning algorithm shows that the lizard's threshold does indeed converge to X^* with probability one. Our proof uses a potential function argument to show that the probability that the foraging threshold remains far from X^* goes to zero. The fact that the learning algorithm converges provides analytic confirmation that the problem of learning optimal foraging thresholds can be solved using biologically plausible quantities. Thus, this work serves to strengthen the conclusions concerning the learning algorithm of Roughgarden [11].

2. Formulation of the Model

In Roughgarden's simulations, prey appear on the ground according to a Poisson point process with intensity a .^{*} In other words, if the ground is divided up into finite patches, and time is divided up into intervals, then the number of prey landing on a given patch of ground during a given time interval is a Poisson random variable whose mean is the product of a , the area of the patch of ground, and the duration of the time interval; furthermore, these various Poisson random variables are mutually independent. Recall that a Poisson random variable with mean λ takes on the value k with probability $e^{-\lambda}\lambda^k/k!$. One basic fact about this model is that with probability 1 no two prey arrive at *exactly* the same instant.

Following Roughgarden, we assume that the lizard surveys an area in front of it shaped like a wedge from a disk, with angle θ (Roughgarden used $\theta = \pi$, making the wedge a semicircle). Thus when the threshold radius is X , prey appear within the surveyed area according to a temporal Poisson process with rate αX^2 , where $\alpha = \theta a/2$. (Since X is a random function of time, the rate αX^2 at which prey appear

^{*} Technically, Roughgarden's actual simulations didn't use continuous time and space, but rather discretized time into units of one second, and discretized space into units of one square meter. But it is clear that this discretization was an artifact of the simulation code rather than an intended feature of the model.

will also be a random function of time.) It is easily seen that when a prey does arrive within the surveyed region, its location is uniformly distributed within the region. In Roughgarden's model, the prey do not move while being chased, so the lizard always catches any prey that it pursues. Again following Roughgarden, we assume that if a prey arrives while the lizard is busy pursuing another prey, the new prey escapes instead of waiting to be caught, and that if a prey arrives outside the lizard's threshold, the lizard does not pursue it, even if it later increases its threshold. Let v be the velocity of the lizard. At time t , let n_t be the number of prey the lizard has captured. For convenience we will assume that n_t increases when the lizard returns to its perch after catching a prey (this assumption will not affect the lizard's behavior). The process starts with some initial threshold X_0 with $n_0 = 1$. (When the process is started, the first prey has just been caught.)

If a lizard decides to pursue a prey that appears at distance d from its perch, then the time taken to chase the prey, catch it, and return to the perch will be $2d/v$. Roughgarden's analysis shows that the foraging threshold predicted by optimal foraging theory is

$$X^* = \left(\frac{3v}{2\alpha} \right)^{1/3}.$$

The learning algorithm proposed by Roughgarden specifies that the lizard will pursue if

$$\frac{t + 2d/v}{n_t + 1} < \frac{t}{n_t}.$$

This expression can be rewritten, to show that the lizard pursues if

$$d < \frac{vt}{2n_t}.$$

Consequently, the learning rule specifies that X_t , the lizard's current guess at the optimal foraging threshold, is $vt/(2n_t)$. Note that X_t increases with time, but decreases every time the lizard returns to its perch after capturing a prey. Let T_i denote the time when the lizard starts looking for the i th prey. If the time that the lizard spends waiting for and pursuing the i th prey is long then $X_{T_{i+1}}$ will be greater than X_{T_i} (so the lizard is unlikely to spend as long waiting for the $(i+1)$ st prey, though it may spend longer pursuing it). However, if the time that the lizard spends waiting for and pursuing the i th prey is short then $X_{T_{i+1}}$ will be smaller than X_{T_i} (so the lizard is likely to wait longer for the $(i+1)$ st prey, though it may spend less time pursuing it). Thus, it is plausible, but not immediately clear, that the threshold values X_t converge over time.

We can describe the distribution of the random variable X_t as follows. Let t_i denote the time at which the i th prey is seen within the lizard's foraging threshold. Let W_i denote the time spent waiting to see the i th prey, d_i denote the distance from the perch to the i th prey, and P_i denote the time spent pursuing the i th prey. We can describe the density functions of d_i , P_i , and W_i (implicitly conditional upon X_{t_i}) as follows.

$$\Pr(d_i \leq y) = \Pr(r \leq (y/X_{t_i})^2) = \Pr(X_{t_i}\sqrt{r} \leq y),$$

where r is a random variable chosen uniformly from the range $[0, 1)$. The first equality follows from the observation that the ratio between the area within radius y and the area within radius X_{t_i} is $(y/X_{t_i})^2$. $P_i = 2d_i/v$, so

$$\Pr(P_i \leq y) = \Pr\left(\frac{2X_{t_i}\sqrt{r}}{v} \leq y\right).$$

We now compute $\Pr(W_i > y)$. Recall that T_i is the time when the lizard starts looking for the i th prey (i.e., $T_i = t_{i-1} + P_{i-1}$). Let the random variable C_y be the number of prey that appear within a radius of $X_{T_i} + vs/2(i-1)$ of the perch at time $T_i + s$, as s goes from 0 to y . (This radius is the threshold radius at time $T_i + s$ provided no prey appear during between times T_i and $T_i + s$.) Note that $\Pr(W_i > y) = \Pr(C_y = 0)$. But C_y is a Poisson random variable, whose parameter is a times the integral of the surveyed area over time. Thus,

$$\Pr(W_i > y) = \exp\left[-\int_0^y \alpha(X_{T_i} + vs/2(i-1))^2 ds\right].$$

Using these definitions, we have

$$\Pr(n_t \geq i) = \Pr\left(\sum_{j=2}^i (W_j + P_j) \leq t\right),$$

so

$$\Pr(X_t \leq y) = \Pr\left(\sum_{j=2}^{\lceil vt/(2y) \rceil} (W_j + P_j) \leq t\right).$$

Let $T(\tau)$ denote the last time at which the lizard has been waiting for τ seconds. (In other words, in addition to a clock measuring real time, there is another clock that

only measures the time τ that the lizard spends sitting at its perch. When the waiting time clock reads τ , the real time clock reads $T(\tau)$. When τ takes on certain special values, of the form $\sum_{i=2}^n W_i$, the real time clock takes on many values due to the positive time required to catch the prey — we break ties by taking the largest such real time, i.e. the time at which the lizard returns to its perch.) Let $m_\tau = n_{T(\tau)}$ and $Y_\tau = X_{T(\tau)}$. Let τ_i be the total amount of time spent waiting before the i th prey is seen within the foraging threshold. The probability that a prey arrives exactly when the lizard returns to the perch after pursuing the i th prey is 0, so with probability 1, $\tau_i < \tau_{i+1}$. In this case, $T(\tau_i)$ is the time at which the lizard returns to its perch after catching the i th prey, so $m_{\tau_i} = i$. Let $X^* = (3v/(2\alpha))^{1/3}$ be the optimal threshold.

3. Overview

Our proof of convergence uses a potential function Φ given by

$$\Phi_\tau = (Y_\tau - X^*)^4 m_\tau^{5/4}.$$

We argue that the expected value of this potential does not get too big, which implies that the threshold cannot deviate much from X^* for too long. We define the potential function in terms of Y_τ and m_τ rather than X_t and n_t because (Y_τ, m_τ) is a Markov process whereas (X_t, n_t) is not, and the Markov property simplifies the analysis.

If at some point the foraging threshold is far from optimal, then it will tend to move towards X^* , causing the expected potential to decrease. If at some point the threshold is exactly X^* , then a short time later it will surely have changed, causing the potential to increase. But the $(Y_\tau - X^*)^4$ term is flat enough when the threshold is near X^* , that this increase in potential will be small. It will turn out that over time the expected increase in potential grows more slowly than the $m_\tau^{5/4}$ term, implying that the deviations from X^* are likely to remain small.

We will argue that regardless of the current threshold radius at a given time, the expected increase in potential in a short interval cannot be too large. Thus, we get a bound on $E[\Phi_{\tau_i}]$. After many prey are found, if Y_{τ_i} deviated from X^* by more than ε_i (with ε_i slowly approaching zero), the potential would be large, so this event cannot be too likely. The potential would need to be large enough to make these events sufficiently unlikely to ensure that they occur for only finitely many i . This constrains the exponent of the m_τ term in the potential to be larger than 1. On the other hand,

if the exponent of the m_τ term were larger than $4/3$, then the expected increase in potential would be too large. Hence, we used $5/4$.

Theorem 1 shows that (with probability 1) $\lim_{i \rightarrow \infty} Y_{\tau_i}$ exists and is X^* , from which it follows that $\lim_{\tau \rightarrow \infty} Y_\tau$ exists, implying that X_t converges to X^* . Our argument uses the expected number of τ_i 's for which Y_{τ_i} deviates far from X^* , since the expected number of τ 's for which Y_τ deviates far from X^* is infinite. The proof of Theorem 1 uses Lemma 5, which shows that the expected value of the potential function at τ_i is bounded by a constant that does not depend on τ .

The proof of Lemma 5 uses Lemmas 3 and 4. Lemma 3 bounds the expected difference between Φ_τ at two times in an interval given the state of the lizard at the first time. Lemma 4 defines a potential Ψ_τ^n that equals Φ_τ in the interval $[\tau_n, \tau_{n+1})$ and equals Φ_{τ_n} and $\Phi_{\tau_{n+1}}$ respectively before and after the interval. Lemma 4 shows that the expected difference between Ψ_τ^n at two times in an interval is bounded by the expected difference between Φ_τ given that n prey have been captured, times the probability that n prey have been captured. Together, Lemmas 3 and 4 are used to show that the expected difference between $\Phi_{\tau_{n+1}}$ and Φ_{τ_n} is less than a slowly decreasing function times the expected waiting time after the n th prey is seen and before the $(n+1)$ st prey is seen within the foraging threshold, from which the result follows.

Lemmas 1 and 2 provide technical details for the proofs of Lemmas 3 and 4. Lemma 3 follows from the definition of Φ_τ . Lemma 4 follows from a case analysis of the number of prey that have been pursued before and during the interval that is being considered.

4. Technical Results

We will use the following technical lemma.

Lemma 1: If P is a Poisson random variable with parameter λ , then

$$\Pr(P \geq k) \leq \lambda^k/k!.$$

Proof:

$$\Pr(P \geq k) = e^{-\lambda} \sum_{i=k}^{\infty} \frac{\lambda^i}{i!} = e^{-\lambda} \frac{\lambda^k}{k!} \sum_{i=0}^{\infty} \frac{\lambda^i k!}{(i+k)!} \leq e^{-\lambda} \frac{\lambda^k}{k!} \sum_{i=0}^{\infty} \frac{\lambda^i}{i!} = \frac{\lambda^k}{k!}.$$

□

Lemma 2: For any fixed finite time interval, there is an upper bound Y such that for all τ in the interval, $Y_\tau \leq Y$. Furthermore, if τ_b and τ_e are drawn from the interval, and $\Delta\tau$ denotes $|\tau_e - \tau_b|$, then the probability that at least k prey arrive within the lizard's foraging threshold between time τ_b and τ_e is at most $O((\Delta\tau)^k)$, where the constant depends only upon the interval.⁺

Proof: The number of prey that appear per unit area per second is a Poisson random variable with parameter a . Thus, the number of prey that appear within the lizard's foraging threshold is dominated by a Poisson random variable with parameter $\alpha Y^2 \Delta\tau$. The lemma follows from Lemma 1. \square

Lemma 3: There is a function $g(n)$ with $g(n) = O(n^{-7/4})$ such that the following is true for any fixed finite time interval: If τ_b and τ_e are drawn from the interval with $\tau_b \leq \tau_e$ and $m_{\tau_b} = n$, then

$$\mathbb{E}[\Phi_{\tau_e} \mid \text{state at time } \tau_b] - \Phi_{\tau_b}$$

is at most

$$(\tau_e - \tau_b)g(n) + O((\tau_e - \tau_b)^2),$$

where the constants in the $O((\tau_e - \tau_b)^2)$ term depend upon n and the interval (in particular upon its maximum foraging radius from Lemma 2) but not upon τ_b or τ_e .

Proof: We start by making a preliminary observation. Suppose that after waiting τ seconds, the lizard sees a prey at time T' within its foraging threshold Y' . By the probability density function of d_{m_τ} , the distance of the prey from the perch is the random variable $Y'\sqrt{r}$, where r is uniformly distributed between 0 and 1. The lizard pursues the prey and returns to its perch. The pursuit time for this prey is $2Y'\sqrt{r}/v$. The new foraging threshold after the prey is caught is

$$Y_\tau = \frac{v(T' + 2Y'\sqrt{r}/v)}{2m_\tau} = \frac{v(2(m_\tau - 1)Y'/v) + 2Y'\sqrt{r}}{2m_\tau} = Y' + Y' \left(\frac{\sqrt{r} - 1}{m_\tau} \right).$$

(Note that the prey will be caught before the time that the lizard has spent waiting, τ , increases.)

⁺ See the appendix for a brief description of the asymptotic notation used in this paper.

Now consider the finite interval in the statement of Lemma 3 and let Y be the interval's foraging threshold bound from Lemma 2. Suppose that τ_b and τ_e are drawn from the interval with $\tau_b \leq \tau_e$ and that $m_{\tau_b} = n$. Let Δ denote $E[\Phi_{\tau_e} \mid \text{state at time } \tau_b] - \Phi_{\tau_b}$. We wish to show that

$$\Delta \leq (\tau_e - \tau_b)g(n) + O((\tau_e - \tau_b)^2),$$

where the constant in the term $O((\tau_e - \tau_b)^2)$ depends upon the interval but not upon τ_b or τ_e and $g(n) = O(n^{-7/4})$ does not depend upon the interval (or upon τ_b or τ_e).

Consider the time period $(\tau_b, \tau_e]$ (the period of time between τ_b and τ_e), and let $\Delta\tau = \tau_e - \tau_b$. Either a prey appears during this period, or not. If no prey appears, the foraging threshold expands by $v\Delta\tau/(2n)$. If exactly one prey appears, then, (by the preliminary observation), the lizard pursues the prey and adjusts the foraging threshold by adding $Z_{\tau_b} = (Y_{\tau_b} + O(\Delta\tau))(\sqrt{r} - 1)/(n + 1)$; meanwhile the foraging threshold expands by $O(\Delta\tau)$. Let $f(y) = (y - X^*)^4$ and let W denote the probability that prey appears during the period. Putting all of this together, we have

$$E[\Phi_{\tau_e} \mid \text{state at time } \tau_b, \text{ choice of } r] \leq (1 - W)n^{5/4}f\left(Y_{\tau_b} + \frac{v\Delta\tau}{2n}\right) + W(n + 1)^{5/4}f(Y_{\tau_b} + Z_{\tau_b} + O(\Delta\tau)) + O((\Delta\tau)^2).$$

So, conditioned on the choice of r ,

$$\begin{aligned} \Delta &\leq n^{5/4}f\left(Y_{\tau_b} + \frac{v\Delta\tau}{2n}\right) - n^{5/4}f(Y_{\tau_b}) \\ &\quad + W\left[(n + 1)^{5/4}f(Y_{\tau_b} + Z_{\tau_b} + O(\Delta\tau)) - n^{5/4}f\left(Y_{\tau_b} + \frac{v\Delta\tau}{2n}\right)\right] + O((\Delta\tau)^2). \end{aligned}$$

We now derive an upper bound for W . Let Y^* denote $Y_{\tau_b} + v\Delta\tau/(2n)$ and note that the lizard's foraging threshold does not exceed Y^* during the period between τ_b and τ_e . The number of prey that appear within the lizard's foraging threshold during this period is dominated by a Poisson random variable with parameter $\alpha\Delta\tau Y^{*2}$. So by Lemma 1, $W \leq \alpha\Delta\tau Y^{*2}$, which is at most $\alpha\Delta\tau(Y_{\tau_b}^2 + Yv\Delta\tau/n + (v\Delta\tau/2n)^2)$. Thus, $W \leq \alpha Y_{\tau_b}^2 \Delta\tau + O((\Delta\tau)^2)$ where the constants in the $O((\Delta\tau)^2)$ depend upon the fixed interval (and therefore, on Y), but not upon τ_b and τ_e . In the interval the second

derivative of f is bounded, so by Taylor's theorem the value of $\Delta/\Delta\tau$ conditioned on the choice of r is as follows:

$$\begin{aligned} \Delta/\Delta\tau \leq n^{5/4} f'(Y_{\tau_b}) \left(\frac{v}{2n} \right) \\ + \alpha Y_{\tau_b}^2 \left[(n+1)^{5/4} f \left(Y_{\tau_b} + Y_{\tau_b} \frac{\sqrt{r}-1}{n+1} \right) - n^{5/4} f(Y_{\tau_b}) \right] + O(\Delta\tau). \end{aligned} \quad (1)$$

Equation 1 led us to choose the exponent $5/4$ in the definition of Φ_{τ_b} . We want the derivative in Equation 1 to be small even when Y_{τ_b} is large. This makes it necessary to have the exponent less than $4/3$. Note that the constants in the $O(\Delta\tau)$ can be chosen uniformly for any τ_b and τ_e in the fixed interval.

Integrating with respect to r , and evaluating Equation 1 at $Y_{\tau_b} = X^*(1+z)$ yields a polynomial in z whose coefficients are functions of n . We need only concern ourselves with the behavior of these functions when n is large enough. Using Maple [10], we find that $\Delta/\Delta\tau$ evaluated at $Y_{\tau_b} = X^*(1+z)$ is equal to the following.

$$\begin{aligned} \Delta/\Delta\tau = \alpha X^{*6} \left[-\frac{1+o(1)}{12} n^{1/4} z^6 - \frac{3+o(1)}{2} n^{1/4} z^5 - \frac{11+o(1)}{4} n^{1/4} z^4 + \frac{11+o(1)}{3} n^{-3/4} z^3 \right. \\ \left. + \frac{1+o(1)}{1} n^{-3/4} z^2 - \frac{2+o(1)}{5} n^{-7/4} z + \frac{1+o(1)}{15} n^{-11/4} \right] + O(\Delta\tau). \end{aligned} \quad (2)$$

where the $o(1)$ terms are functions of n converging to 0 as $n \rightarrow \infty$ (independently of z).

Suppose that n is sufficiently large. Ignoring the $O(\Delta\tau)$ term for the moment, let β denote the first term. We consider the following cases.

Case 1

$$z \geq 3n^{-1/2}$$

In this case, the term $[(11+o(1))/4]n^{1/4}z^4$ is at least three times as large as each of the three positive terms, so $\beta \leq 0$.

Case 2

$$-1 \leq z \leq -3n^{-1/2}$$

In this case, the term $[(11+o(1))/4]n^{1/4}z^4$ is at least $3/2$ times as large as the term $[(3+o(1))/2]n^{1/4}z^5$ and at least nine times as large as each of the three other positive terms. Thus, $\beta \leq 0$.

Case 3

$$|z| \leq 3n^{-1/2}$$

In this case, each term is $O(n^{-7/4})$, so $\beta = O(n^{-7/4})$.

Note that $z \geq -1$. Thus we have that $\Delta/\Delta\tau = O(n^{-7/4}) + O(\Delta\tau)$ for all values of z , for all values of τ_b and τ_e contained in a given finite interval. \square

We will find it convenient to define

$$\Psi_\tau^n = \begin{cases} \Phi_{\tau_n}, & \text{if } \tau < \tau_n; \\ \Phi_\tau, & \text{if } \tau_n \leq \tau < \tau_{n+1}; \\ \Phi_{\tau_{n+1}}, & \text{if } \tau_{n+1} \leq \tau. \end{cases}$$

As a function of τ , Ψ_τ^n is discontinuous at τ_{n+1} but it is continuous elsewhere.

Lemma 4: For any fixed finite time interval, if τ_b and τ_e are drawn from that interval then

$$\mathbb{E} [\Psi_{\tau_e}^n - \Psi_{\tau_b}^n] = \mathbb{E} [\Phi_{\tau_e} - \Phi_{\tau_b} \mid m_{\tau_b} = n] \cdot \Pr(m_{\tau_b} = n) + O((\Delta\tau)^2),$$

where the constants in the $O((\Delta\tau)^2)$ term depend upon n and the fixed interval's maximum foraging radius from Lemma 2.

Proof: Let Y be the interval's maximum foraging radius from Lemma 2 and let $\gamma_r = \mathbb{E} [\Psi_{\tau_e}^n - \Psi_{\tau_b}^n \mid m_{\tau_b} = r]$. By the linearity of expectation,

$$\mathbb{E} [\Psi_{\tau_e}^n - \Psi_{\tau_b}^n] = \sum_{r=1}^{\infty} \gamma_r \Pr(m_{\tau_b} = r). \quad (3)$$

We now consider possible values of r .

Case 1: $r > n$: In this case, $\gamma_r = 0$.

Case 2: $r \leq n - 1$: In this case, if no prey are seen then $\Psi_{\tau_e}^n - \Psi_{\tau_b}^n = 0$. If one prey is seen (by Lemma 2, this happens with probability $O(\Delta\tau)$), then $|\Psi_{\tau_e}^n - \Psi_{\tau_b}^n| = O(\Delta\tau)$. Even if two or more prey are seen (by Lemma 2, this happens with probability $O((\Delta\tau)^2)$), $|\Psi_{\tau_e}^n - \Psi_{\tau_b}^n|$ is still $O(1)$. Hence $\gamma_r = O((\Delta\tau)^2)$.

Case 3: $r = n$: In this case, let $\Delta\Phi = \Phi_{\tau_e} - \Phi_{\tau_b}$ and $\Delta\Psi = \Psi_{\tau_e}^n - \Psi_{\tau_b}^n$. If k prey are found, then $|\Delta\Phi| \leq Y^4(n+k)^{5/4}$ and $|\Delta\Psi| \leq Y^4(n+1)^{5/4}$, so

$$|\Delta\Psi - \Delta\Phi| \leq 2Y^4(n+k)^{5/4}.$$

If $k = 0$, then $\Delta\Psi = \Delta\Phi$. If $k = 1$ (with probability $O(\Delta\tau)$, by Lemma 2) then $|\Delta\Psi - \Delta\Phi| = O(\Delta\tau)$. Using Lemmas 1 and 2,

$$\begin{aligned} |E[\Delta\Psi \mid m_{\tau_b} = n] - E[\Delta\Phi \mid m_{\tau_b} = n]| &\leq O((\Delta\tau)^2) + 2Y^4 \sum_{k=2}^{\infty} (n+k)^{5/4} (\alpha Y^2 \Delta\tau)^k / k! \\ &= O((\Delta\tau)^2). \end{aligned}$$

Since $\gamma_r = \mathbb{E}[\Delta\Psi \mid m_{\tau_b} = r]$,

$$\gamma_r = \mathbb{E} \left[\Phi_{\tau_e} - \Phi_{\tau_b} \mid m_{\tau_b} = r \right] + O((\Delta\tau)^2).$$

Summing over r according to Equation 3, we get

$$\mathbb{E} \left[\Psi_{\tau_e}^n - \Psi_{\tau_b}^n \right] = \mathbb{E} \left[\Phi_{\tau_e} - \Phi_{\tau_b} \mid m_{\tau_b} = n \right] \cdot \Pr(m_{\tau_b} = n) + O((\Delta\tau)^2),$$

where the constants in the $O((\Delta\tau)^2)$ term depend upon n and upon the fixed interval.

□

Lemma 5: For all n , $\mathbb{E}[\Phi_{\tau_n}] \leq \Phi_0 + O(1)$. (The $O(1)$ term does not depend on τ .)

Proof: We start by computing $\Phi_{\tau_{n+1}} - \Phi_{\tau_n} = \Psi_{\infty}^n - \Psi_0^n$. To do this, we will use Lemmas 3 and 4. Both lemmas may be applied to any finite interval, so we will apply them to the intervals $[0, 1], [1, 2], \dots$. Any interval $[i, i + 1]$ may be subdivided into subintervals of size $1/M$ (for any $M > 1$). Applying lemmas 3 and 4 to each subinterval and summing gives

$$\begin{aligned} \mathbb{E} \left[\Psi_{i+1}^n - \Psi_i^n \right] &= \sum_{j=1}^M \mathbb{E} \left[\Psi_{i+j/M}^n - \Psi_{i+(j-1)/M}^n \right] \\ &\leq \sum_{j=1}^M \left[(g(n)/M + O(1/M^2)) \Pr(m_{i+(j-1)/M} = n) + O(1/M^2) \right], \end{aligned}$$

where the constants in the $O(1/M^2)$ terms depend on the interval $[i, i + 1]$ and on n but not the subintervals. Taking the limit of both sides as $M \rightarrow \infty$ gives

$$\mathbb{E} \left[\Psi_{i+1}^n - \Psi_i^n \right] \leq g(n) \int_i^{i+1} \Pr(m_\tau = n) d\tau.$$

Summing over all i we get

$$\mathbb{E}[\Phi_{\tau_{n+1}} - \Phi_{\tau_n}] \leq g(n) \int_0^\infty \Pr(m_\tau = n) d\tau. \quad (4)$$

The integral in Equation 4 is the expected waiting time after the n th prey is seen before the $(n + 1)$ st prey is seen. We will show that this is $O(n^{2/3})$. To see this, consider the (unlikely) event that the lizard has spent $n^{2/3}$ seconds waiting for the $(n + 1)$ st prey,

and hasn't seen it yet. Then the total waiting time must be at least $n^{2/3}$ seconds, and the foraging threshold, $\frac{vt}{2n_t}$, is then at least

$$\frac{vn^{2/3}}{2n} = \frac{v}{2}n^{-1/3}.$$

So the area of the region that the lizard surveys from this time until the $(n+1)$ st prey is at least $(\theta v^2/8)n^{-2/3}$. The expected time until a prey appears within just this patch of ground is $(8/\theta v^2)n^{2/3}/a$. Thus, the expected time spent waiting for the $(n+1)$ st prey is at most

$$n^{2/3} + \frac{8}{a\theta v^2}n^{2/3} = O(n^{2/3}).$$

Since $g(n) = O(n^{-7/4})$ (from Lemma 3), we conclude that $E[\Phi_{\tau_{n+1}} - \Phi_{\tau_n}] = O(n^{-13/12})$. Since $n^{-13/12}$ is a convergent series, the expected value of the potential just after the n th prey is caught is at most Φ_1 plus some constant independent of n . \square

Theorem 1: *With probability one $X_t \rightarrow X^*$.*

Proof: First note that if the lizard returns to the perch at time t after catching the n th prey and the lizard returns to the perch at time t' after catching the $(n+1)$ st prey, then the foraging threshold during the time interval $[t, t']$ stays in the interval $[vt/(2(n+1)), vt'/(2n)]$. Equivalently, the foraging threshold during the time interval $[t, t']$ stays in the range $[(vt/(2n))(1 - 1/(n+1)), (vt'/(2(n+1)))(1 + 1/n)]$. Note that for all j , with probability 1, Y_{τ_j} is the foraging threshold just after the j th prey is caught. Thus, $Y_{\tau_n} = vt/(2n)$ and $Y_{\tau_{n+1}} = vt'/(2(n+1))$. Thus, if n is sufficiently large, then the foraging threshold during the time interval $[t, t']$ stays in the interval $[X^* - 1/n^{1/33}, X^* + 1/n^{1/33}]$ unless either $Y_{\tau_n} \notin [X^* - 1/n^{1/32}, X^* + 1/n^{1/32}]$ or $Y_{\tau_{n+1}} \notin [X^* - 1/(n+1)^{1/32}, X^* + 1/(n+1)^{1/32}]$.

Therefore we focus on the subsequence of foraging thresholds Y_{τ_n} just after the n th prey was caught.

Using Markov's inequality and recalling that, with probability 1, $m_{\tau_n} = n$, we get

$$\begin{aligned} \Pr(|X^* - Y_{\tau_n}| > \varepsilon) &= \Pr(n^{5/4} |X^* - Y_{\tau_n}|^4 > n^{5/4} \varepsilon^4) \\ &= \Pr(\Phi_{\tau_n} > n^{5/4} \varepsilon^4) \\ &\leq E[\Phi_{\tau_n}] n^{-5/4} \varepsilon^{-4}. \end{aligned}$$

If we take $\varepsilon = n^{-1/32}$ and note that, by Lemma 5, $E[\Phi_{\tau_n}]$ is bounded from above by a constant, then we get

$$\Pr(|X^* - Y_{\tau_n}| > n^{-1/32}) = O(n^{-9/8}).$$

Thus, as the process runs, the expected number of integers n such that the Y_{τ_n} deviates from X^* by more than $n^{-1/32}$ is $O(1)$. Almost surely, there are a finite number of such integers n , so with probability one, there are only a finite number of integers n such that during the time between the catching of the n th prey and the $(n + 1)$ st prey, the foraging threshold goes outside of the range $[X^* - 1/n^{1/33}, X^* + 1/n^{1/33}]$. \square

5. Discussion

We have analyzed a simplified version of the general learning problem that faces the anole. Using very weak assumptions on the distribution of prey, we are able to prove that the sequence of foraging thresholds used by the lizard converges to the foraging threshold predicted by optimal foraging theory. This analysis confirms the conclusions made by Roughgarden [11] concerning the ability of this learning rule to optimize a lizard's foraging behavior.

Roughgarden [11] describes a more general version of this problem that includes more details concerning the energy expenditures of the lizard. Specifically, Roughgarden assumes that a lizard expends e_w energy per time while waiting and e_p energy per time while pursuing a prey. The energy obtained from a prey is e . An optimal threshold similar to X^* can be derived for this model. An important difference of this model is that the lizard's behavior aims to maximize the current energy gained from a prey. The model that we have analyzed assumes that $e_p = e_w$, in which case minimizing the time is equivalent to maximizing the energy. For other parameters, this equivalence does not apply and a different learning rule and convergence analysis is needed.

This limits the generality of the result that we have presented. The type of "sit-and-wait" predation model that we have analyzed may be applicable to other organisms. Minimizing the average time spent per prey captured may be reasonable model of optimal behavior in many contexts, and the the main challenge for adapting our analysis to such organisms would be to account for the specific distribution of prey in the organism's environment. For organisms for which this model is not appropriate, a more general model described by Roughgarden [11] would need to be analyzed.

Two other generalizations considered by Roughgarden model are (i) the ability of prey to escape from the lizard and (ii) the existence of different categories of prey that have differing energetic value. From Roughgarden’s discussion of model (ii), it is clear that a convergence proof for the previous model would immediately generalize to this model. However, for model (i) a convergence analysis will be much more difficult. In fact, Roughgarden is unable to determine an analytic value for the foraging threshold predicted by optimal foraging for a specific model of prey behavior.

One assumption made by these models is that the lizard does not compete with other lizards while learning its optimal foraging threshold. While Roughgarden has considered models that allow this competition [12], his justification for the current model is that the foraging threshold simply determines how many lizards there can be in a region.

Finally, we note that other learning rules might yield simpler analysis. Erik Ordentlich [9] has observed that the time that the lizard would take to run to the threshold and back is equal to the average time that the lizard has spent per prey. We suspect that other learning rules based on this observation might also converge. Nick Littlestone has observed [5] that the time that the lizard spends waiting is, in the limit, half of the time that the lizard spends running. This suggests a simple learning rule – adjust the radius whenever the ratio between the running time and waiting time is wrong. This learning rule only requires division by two and not general division, which may be more biologically plausible.

Appendix

In this paper, we use asymptotic notation to specify bounds on how functions behave in the limit. For details about this notation see (for example) Chapter 2 of [1]. To aid the reader, we give brief definitions of the notation (from [1]) here.

For functions f and g , we say that $f(n) = o(g(n))$ if and only if $\lim_{n \rightarrow \infty} \frac{f(n)}{g(n)} = 0$. We say that $f(n) = O(g(n))$ if and only if there are positive constants c and n_0 such that for all $n \geq n_0$ we have $f(n) \leq cg(n)$.

When we say that a quantity is “at most $O(g(n))$ ”, we mean that there is a function $f(n)$ such that $f(n) = O(g(n))$ and the quantity is at most $f(n)$. Similarly, when we say “ $f(n) = g(n) + O(h(n))$ ”, we mean that there is a function $j(n) = O(h(n))$ such that $f(n) = g(n) + j(n)$.

References

- 1 T.H. CORMEN, C.E. LEISERSON AND R.L. RIVEST, *Introduction to Algorithms* (MIT Press 1990).
- 2 W. FELLER, *An Introduction to Probability Theory and its Applications*, Volume 1, 3rd Edition, John Wiley and Sons, 1968.
- 3 J. KOZA, *Genetic programming: On Programming Computers by Means of Natural Selection and Genetics*, The MIT Press, Cambridge, MA, 1992, pp. 314–328.
- 4 J. R. KREBS AND N. B. DAVIES, *Behavioral Ecology. Second Edition*. Blackwell, Oxford, 1984.
- 5 N. LITTLESTONE, *Personal communication*, July, 1996.
- 6 M. MANGEL AND C. W. CLARK, *Dynamic Modeling in Behavioral Ecology*. Princeton University Press, Princeton N.J., 1988.
- 7 J. M. MCNAMARA, An optimal sequential policy for controlling a Markov renewal process, *J. Appl. Prob.*, 22, 324-335, 1985.
- 8 F. MENCZER, W. E. HART, AND M. L. LITTMAN, Appendix to S. Shafir and J. Roughgarden's 'The Effect of Memory on Individual Fitness in a Lizard', *Adaptive Individuals in Evolving Populations: Models and Algorithms*, Ed. R. K. Belew and M. Mitchell, Addison-Wesley, SFI Studies in the Sciences of Complexity Vol. 26, 183-186, 1996.
- 9 E. ORDENTLICH, *Personal communication*, July, 1996.
- 10 D. REDFERN, *The Maple Handbook*, Springer-Verlag, New York, 1993.
- 11 J. ROUGHGARDEN, *Anolis lizards of the Caribbean: Ecology, evolution and plate tectonics*. Oxford Univ. Press, 1995.
- 12 J. ROUGHGARDEN, *Personal communication*, September, 1995.
- 13 S. SHAFIR AND J. ROUGHGARDEN, The effect of memory length on individual fitness in a lizard, *Adaptive Individuals in Evolving Populations: Models and Algorithms*, Ed. R. K. Belew and M. Mitchell, Addison-Wesley, SFI Studies in the Sciences of Complexity Vol. 26, 173-182, 1996.