

Report

Reinforcement Learning Enables Resource Partitioning in Foraging Bats

Aya Goldshtein,¹ Michal Handel,¹ Ofri Eitan,¹ Afrine Bonstein,¹ Talia Shaler,¹ Simon Collet,² Stefan Greif,³ Rodrigo A. Medellín,⁴ Yuval Emek,⁵ Amos Korman,^{2,*} and Yossi Yovel^{1,3,6,*}

¹School of Zoology, Faculty of Life Sciences, Tel Aviv University, Tel Aviv 6997801, Israel

²The Research Institute on the Foundations of Computer Science (IRIF), CNRS and University of Paris, Paris 75013, France

³Sagol School of Neuroscience, Tel Aviv University, Tel Aviv 6997801, Israel

⁴Departamento de Ecología de la Biodiversidad, Instituto de Ecología, Universidad Nacional Autónoma de México, Ciudad de México 04510, México

⁵Faculty of Industrial Engineering and Management, Technion - Israel Institute of Technology, Haifa 3200003, Israel

⁶Lead Contact

*Correspondence: amos.korman@irif.fr (A.K.), yossiyovel@gmail.com (Y.Y.)

<https://doi.org/10.1016/j.cub.2020.07.079>

Konstanzer Online-Publikations-System (KOPS)

URL: <http://nbn-resolving.de/urn:nbn:de:bsz:352-2-kcw6lfw3e92k4>

SUMMARY

Every evening, from late spring to mid-summer, tens of thousands of hungry lactating female lesser long-nosed bats (*Leptonycteris yerbabuenae*) emerge from their roost and navigate over the Sonoran Desert, seeking for nectar and pollen [1, 2]. The bats roost in a huge maternal colony that is far from the foraging grounds but allows their pups to thermoregulate [3] while the mothers are foraging. Thus, the mothers have to fly tens of kilometers to the foraging sites—fields with thousands of Saguaro cacti [4, 5]. Once at the field, they must compete with many other bats over the same flowering cacti. Several solutions have been suggested for this classical foraging task of exploiting a resource composed of many renewable food sources whose locations are fixed. Some animals randomly visit the food sources [6], and some actively defend a restricted foraging territory [7–11] or use simple forms of learning, such as “win-stay lose-switch” strategy [12]. Many species have been suggested to follow a trapline, that is, to revisit the food sources in a repeating ordered manner [13–22]. We thus hypothesized that lesser long-nosed bats would visit cacti in a sequenced manner. Using miniature GPS devices, aerial imaging, and video recordings, we tracked the full movement of the bats and all of their visits to their natural food sources. Based on real data and evolutionary simulations, we argue that the bats use a reinforcement learning strategy that requires minimal memory to create small, non-overlapping cacti-cores and exploit nectar efficiently, without social communication.

RESULTS AND DISCUSSION

The foraging task of the lesser long-nosed bat is particularly complex. Saguaro fields (*Carnegiea gigantea*) in the northern Sonora contain many hundreds of cacti per square kilometer [23]. The cactus flowers open in the evening and start producing nectar ca. 1.5 h after sunset at a constant rate of 0.1 ± 0.004 mL/h per flower (Figure S1A; STAR Methods). The quality of a specific cactus can vary dramatically according to the number of flowers it opens on a given night, which translates linearly to its total rate of nectar production (flowers produce nectar at similar rates and sugar concentration is similar; STAR Methods). Moreover, there is little correlation between the quality of neighboring cacti (Figure 1A).

Energetic estimates (STAR Methods) suggest that a lactating female needs to drink 14.1–18.9 mL of nectar per night. Because a single cactus provides no more than 0.65 mL during the hours that a bat spends in the field, a bat must visit many cacti. As high-quality cacti are rare and most cacti have 0 to 1 open flowers (Figures 1A and S1B), it is worthwhile for a bat to invest time in exploring the field. Moreover, bats have to re-map the field every night. Although we found a significant correlation between the

number of flowers on a cactus on consecutive nights, there was much inter-night variability (Spearman correlation test: $r = 0.6$; $p < 0.001$; Figures S1C and S1D). Importantly, even if the exploration is effective, the corresponding exploitation might not be trivial. In theory, a bat could memorize when it last visited a profitable cactus [24] and plan to revisit it after some time interval, but it cannot be certain it will find the expected nectar, because a conspecific might deplete it in the meanwhile (the bats' behavior suggests that they do not know whether a flower was recently visited; Figures S1E and S1F). We thus set out to examine how bats make foraging decisions under such uncertainty.

We used miniature GPS devices to track bats' movement. We employed drone and aerial imaging (of ~3,600 cacti) to identify all resources available to the bats, that is, the positions of the cacti that they visited and the estimated distribution of flowers. To complete our understanding at the population level, we video-monitored tens of cacti over full nights, recording all bat (and other pollinator) activity. Finally, all of these data were used to develop a mathematical model that aims to explain the bats' foraging strategy.

In total, we GPS tracked 17 lesser long-nosed females. We then used aerial imaging to reconstruct the location of all cacti



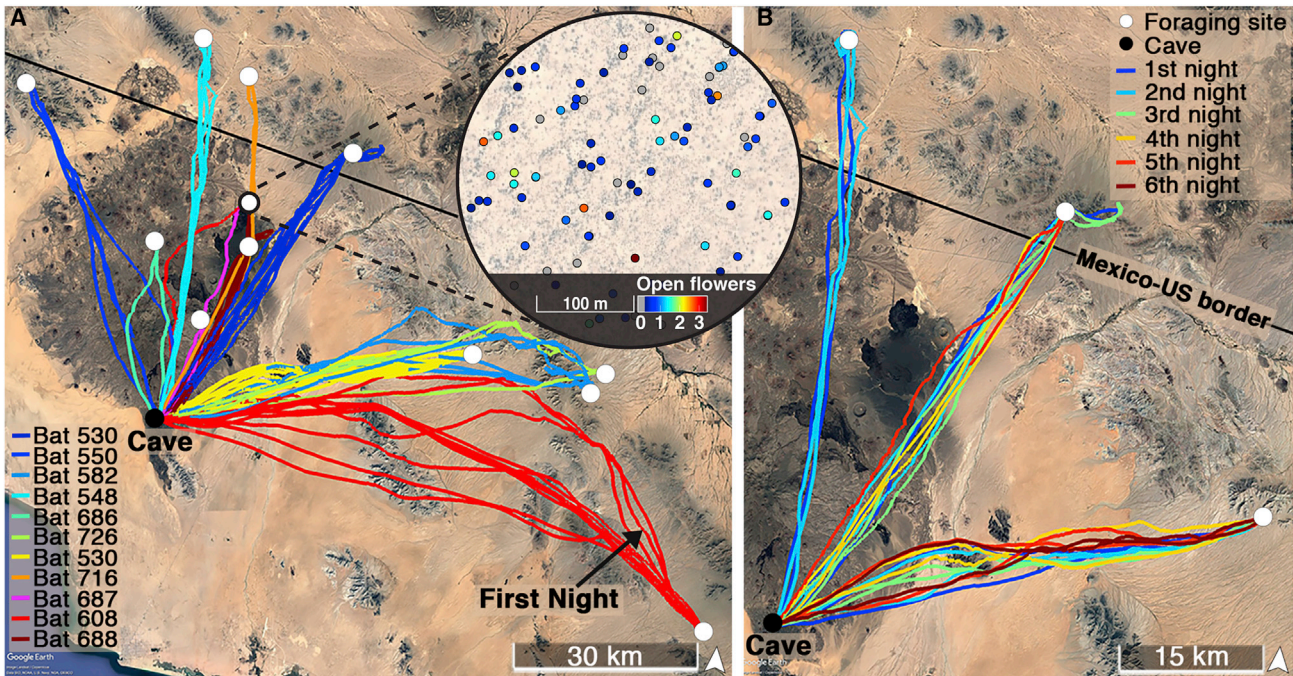


Figure 1. Typical Movement of Lactating Nectar-Feeding Bats

(A) Bats fly tens of kilometers from the cave to their foraging site (colors represent different individuals; $n = 11$ bats for which we had at least one full night). The bat that flew farthest (red) flew a one-way distance of ~ 104 km on her first night (depicted by a black arrow) to her foraging site (white circle). The zoomed-in circle panel shows how the cacti quality—the number of open flowers—differ greatly within one of the Saguaro fields, with 0–4 open flowers on an average night (0 flowers are presented as gray circles and 1–4 flowers are presented as blue to red circles). See also [Figure S1](#).

(B) Bats typically visited the same Saguaro field on consecutive nights (colors from blue to red depict different nights for each bat; three bats with a mean of 4.3 ± 2.0 nights per bat are shown). When returning to the same site, the bats visited $\sim 75\%$ of the cacti that they visited on the previous night ($n = 8$ bats; see also [STAR Methods](#) and [Table S1](#)).

at the foraging sites of eight bats ([STAR Methods](#)). Henceforward, unless stated otherwise, all foraging analyses were conducted for these $n = 8$ bats, with a mean of 2.9 ± 2.1 nights per bat. Tracking revealed extreme commutes of tens of kilometers, from the colony to the cacti fields (the average one-way commute was 55.4 ± 17.2 km, mean \pm SD, $n = 17$ bats, with a record-holding bat that flew ~ 104 km to reach the field; [Figure 1A](#)). On consecutive nights, bats typically returned to the same field (in 25 out of the 28 consecutive nights that we monitored for 8 bats; [Figure 1B](#)).

Commuting to the foraging fields took 1.6 ± 0.8 h, in which bats flew at a mean ground speed of 8.7 ± 0.8 m/s, mean air speed of 9.1 ± 0.9 m/s, and a mean altitude of 126 ± 75 m above ground (mean \pm SD; $n = 17$ bats; maximum ground speed was 18.3 m/s and maximum altitude was $\sim 1,045$ m). These long commute flights account for an average of $\sim 46\%$ of the total time a bat spends outdoors, implying that the bats should use the remaining time for exploiting the field efficiently.

Foraging Dynamics throughout the Night

Once in the field, bats spent 2.2 ± 0.7 h foraging ($n = 11$; mean \pm SD). Even though a single Saguaro field in this region can spread over a few square kilometers, each bat concentrated its foraging on a relatively small area (0.14 ± 0.09 km²; the convex hull of all locations defined as foraging; [STAR Methods](#)). Moreover, even within this restricted site, the bat did not visit all cacti uniformly

but clearly preferred some cacti over others ([Figure 2A](#)). We defined a cactus that was visited by a specific bat at least five times throughout the night as belonging to this bat's core cacti (see [STAR Methods](#) and [Figures S1G](#) and [S1H](#) for the reasoning behind this definition).

The bats combined exploitation of their core cacti with exploration of new unvisited cacti ([Figures 2B](#) and [2C](#) and inset in [Figure 2A](#)), and they gradually increased the exploitation-to-exploration ratio at a steady rate throughout the night ([Figure S1I](#)). By the end of the night, a bat visited an average of 86 ± 48 cacti, out of which 28 ± 26 ($\sim 33\%$) cacti were part of its core cacti (mean \pm SD; [Figures 2B](#), [2C](#), and [S1J](#)). This accounts for drinking ~ 18 mL of nectar per night, approximately the amount required by a lactating female ([STAR Methods](#)).

Information-Based Foraging

Bats clearly did not visit cacti in the field in a deterministic ordered sequence, as a trapline strategy would suggest (see [Figures S1K](#) and [S1L](#) for more details). However, the bats also did not visit the cacti at random. Indeed, video analysis revealed that cacti with more flowers were visited significantly more often (Spearman correlation test; $r = 0.40$; $p = 0.02$; $n = 33$ cacti; [Figure S1M](#)), suggesting that some kind of a reinforcement mechanism underlies the bats' behavior.

In order to unravel the bats' visitation strategy, we developed a mathematical model that simulates a single night of foraging

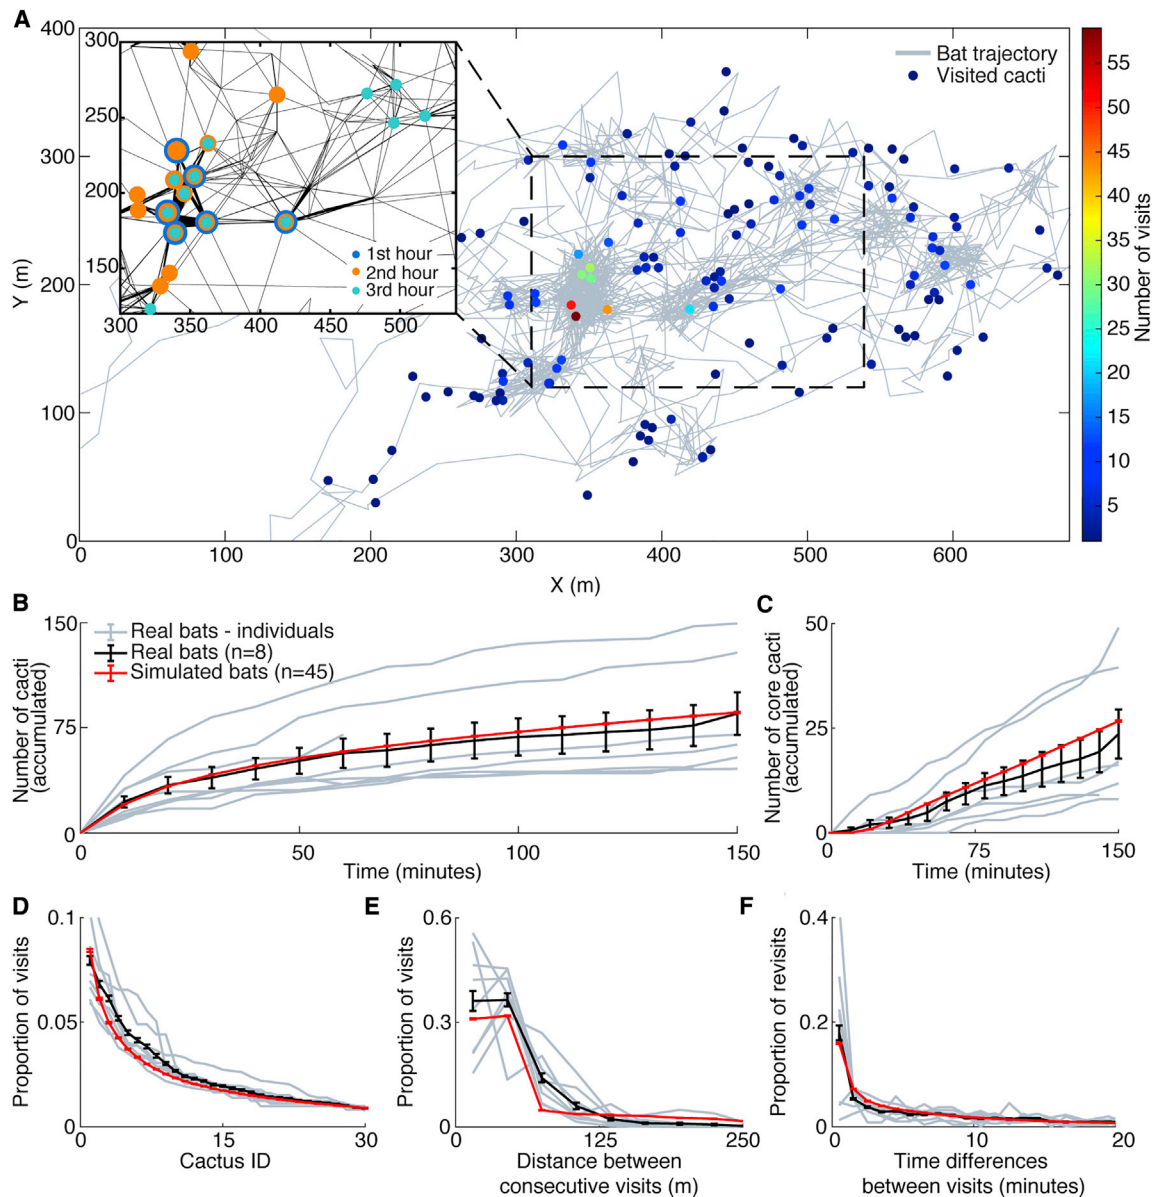


Figure 2. Nectar-Feeding Bats Foraging Dynamics

(A) Flight trajectory of one bat is shown (gray line), and the cacti it visited are colored according to the number of visits (unvisited cacti are not presented). Top left inset shows hops between cacti, where the width of the black lines represents the number of hops between cacti and colored circles represent only the core cacti that were visited more than 5 times at different hours throughout the night. In (B)–(F), colors are as follows: gray, real bats data; black, mean \pm SE of 8 real bats; red, mean \pm SE of 45 simulated bats (30 simulations).

(B and C) The accumulated number of (B) visited cacti and (C) core cacti throughout the night (see also Figure S1J).

(D) The proportion of visits per cactus are sorted according to cactus ID, i.e., cactus 1 is the cactus that received most visits.

(E) The distribution of distances between consecutive visited cacti.

(F) The distribution of time lags between revisits to the same cactus. Core cacti were close to each other (25.1 ± 13.6 m apart), and bats revisited them very often (consecutive visits to core cacti were on average 8.9 ± 3.5 min apart; E and F).

under competition in a space with multiple replenishing food sources, using a rule inspired by reinforcement learning [25]. In the simulations, bats start without prior knowledge on cacti quality and remember the value (“nectar weight”) and positions of the ten best cacti they experienced (Figure S1J). At each visit, the bat updates the nectar weight of the cactus

depending on the cactus’ nectar yield as experienced by the bat in previous visits and according to a learning rate α . Then, the bat samples its next visit from the set of cacti that includes the ten remembered best cacti and the other cacti around it (STAR Methods). The probability to pick a cactus in this set is proportional to its nectar weight, i.e., to the amount

of nectar it provides, and inverse proportional to its distance from the current location of the bat.

Importantly, we did not fit the data in order to find the learning rate α . We ran evolutionary simulations, in which multiple bats with different individual learning rates competed with each other. Bats that accumulated more nectar had higher chances to reproduce, and the learning rates in each generation were slightly mutated. After many generations, the simulated bats converged to a learning rate of 4.6 (STAR Methods). The simulated bats with this α behaved very similarly to the real bats (compare the red and black lines in Figures 2B–2F).

Similar to the real bats, the simulated bats showed a strong preference for specific core cacti, but they also continued exploring other cacti (Figures 2B and 2C). The simulated bats increased the exploitation-to-exploration ratio in a similar rate as the real bats, and they showed a similar preference for cacti with more flowers (Figures S1I and S1M, respectively). The simulated bats hopped between cacti with a similar distance distribution as the real bats (Figure 2E), and they revisited cacti after similar time intervals (Figure 2F). Our model was also able to predict the nectar consumption per visit of the actual bats, as we observed in the video (Figure S1N). All the above suggests that our simple model, which assumes little memory and a single free parameter, can well approximate the bats' behavior.

We also compared our model to two other models. The first had no learning, and bats visited the cacti uniformly at random. The second model, a version of a trapline model, is similar to our model, except that cacti at the core are visited in an ordered manner that minimizes the distance to the next cactus (Figures 3A–3E; STAR Methods). Both of these alternative models performed worse in terms of their fit to the actual data (Figure 3F) and were also significantly less efficient in terms of nectar consumption (one-way ANOVA: $F(2) = 28.5$; $p < 0.001$; Figure 3G).

We conclude that learning is a key feature of the bats' behavior. Because bats rarely encounter conspecifics at the cacti (see below), the main way in which a bat experiences competition is through depletion of cacti. We thus hypothesized that learning is particularly useful in the face of competition. To examine this, we ran another simulation comparing the model, in which cacti are visited uniformly at random in each step, to our reinforcement algorithm strategy under two circumstances: (1) when the forager is alone in the field and (2) when it forages with conspecifics—unlike the simulations above, here, we tested a situation where half of the bats use the first strategy (i.e., random) and the other half use the second strategy (i.e., learning). In the case of a sole forager, a purely random model was as good as the learning model (permutation-based t test: $p = 0.3$; $n = 10,000$ simulations). This phenomenon was reversed in the case of many competing bats, where the simulated bats that used reinforcement learning consumed 43% more nectar than the randomly visiting bats (17.2 mL versus 11.9 mL in the learning and random models, respectively; permutation-based t test: $p < 0.0001$; $n = 10,000$).

Individual Foraging Cores with Little Overlap Emerge from Reinforcement Learning

A prominent feature of our model is the formation of foraging cores with little overlap between individuals (Figure 3H; we prefer the term “core” to “home range” [26] because their important

feature is the set of cacti they are composed of and not their area). The sharing of core cacti by multiple bats in our simulations was significantly lower than expected if the bats randomly visited the cacti (STAR Methods; permutation-based t test: $p = 0.0001$; $n = 10,000$ simulations). Data suggest that, in reality, bats also form disjoint cores with little overlap—video analysis revealed an estimated overlap of 1.3 bats per cactus, although in the simulations, it was 1.5 bats per cactus (Figures S2A and S2B).

Importantly, foraging cores are unlikely to be a result of territorial behavior: bats' core cacti spread over 0.04 km², an area far too large for a bat to defend—flying from one side to the other would require ~40 s although a visit of a conspecific to one of the bat's cacti requires less than 1 s (Video S1). Moreover, the range from which a bat could detect a conspecific based on its echolocation is limited [27] and visual detection range is even more limited, reaching only ~50 m [27, 28].

To further understand the potential impact of interference competition, we added to the simulation a component of aggressive defense where a bat deters conspecifics it encounters at its core cacti (Figures 3A–3E, blue lines; STAR Methods). Consistent with our hypothesis that core formation is not a result of territorial defense, this aggressive model did not differ from the previous non-aggressive model.

Moreover, the formation of foraging cores was also not a result of the limited memory of the simulated bats (10 remembered cacti). We ran the model with unlimited memory, where the bats remember the positions and weights of all cacti in the field, and the same type of core cacti evolved (Figures S2C–S2F). In fact, what appears to determine the size of a foraging core is the learning rate α , where a larger α generates smaller cores (Figures S2G–S2J). This is intuitive, as a larger α drives more exploitation of previous profitable cacti at the expense of exploring new cacti.

The little overlap between neighboring cores in the simulation appears to result from a combination of two factors. First, bats initially visit cacti randomly, and then a positive feedback effect enhances the profitable initial discoveries through reinforcement learning. Second, the nectar weights of profitable cacti are reinforced accumulatively, increasing their weight in the eyes of the owner bat and, at the same time, decreasing their nectar weights in the eyes of occasional visitors that are likely to find the cacti empty. When this reinforcement learning strategy is employed in parallel by many bats, it enables the partitioning of the resources in the field without active communication and reduces the number of conflicts. Indeed, very few social agonistics interactions were observed—only in 5% of the 1,123 visits documented on video.

The Efficiency of Core Formation

Despite its simplicity, the proposed reinforcement learning strategy results in highly efficient foraging: as a group, the simulated bats consumed 91% of the total amount of nectar that was accumulated in the field during the night (14.9 ± 0.8 mL per bat). This means, in particular, that any other foraging algorithm could not improve nectar consumption by more than 10% on average.

In principle, one might expect that a more sophisticated (non-realistic) algorithm, that, e.g., maintains a complete memory of the qualities of all cacti and visits them accordingly, would

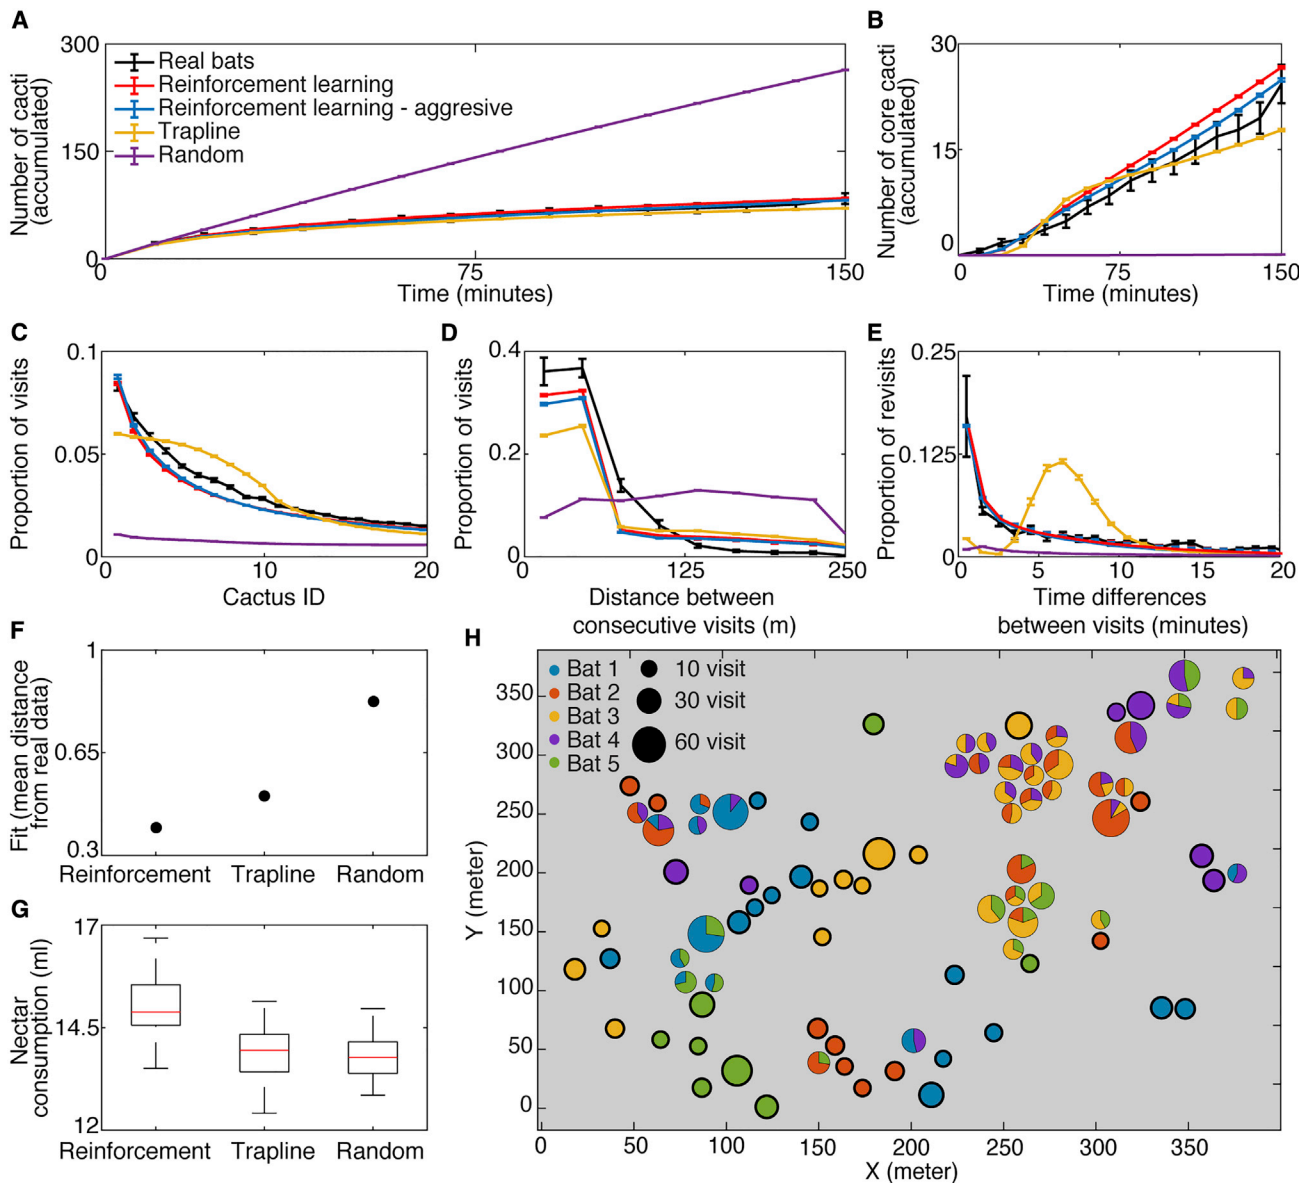


Figure 3. Comparison of Different Models

(A–E) Foraging behavior of real bats (black, mean \pm SE, $n = 8$ bats) and simulated bats (mean \pm SE, $n = 45$ bats) using different foraging models: reinforcement learning (red); reinforcement learning with aggression (blue); trapline (yellow); and random (purple).

(A and B) Accumulated number of (A) visited cacti and (B) core cacti (first 2.5 h).

(C) Sorted proportion of visits per cactus (in the 20 most visited cacti).

(D) The distribution of distances between consecutive visited cacti (30-m bins).

(E) The distribution of time lags between revisits to the same cactus (first 20 min).

(F) The fit of different foraging models (normalized mean distance from real data; data were normalized to a max of 1 before fitting).

(G) Bats' nectar consumption in different foraging models.

(H) Core cacti of five simulated bats in the reinforcement learning model. Cacti of different individuals are depicted by different colors, and circle size represents the number of visits (accumulated over a full night). Cacti that were in the core of more than one bat appear in multiple colors according to the relative number of visits of each bat, and cacti that were in the core of a single bat are circled by a black line.

See also [Figure S2](#).

perform significantly better than one that only selects a small subset of core cacti. However, an analytic analysis of a simplified model suggests that this intuition is wrong (see [29]). We mathematically analyzed a simplified setting in which a (single) bat aims

to exploit a set of n cacti with different refill rates (r_i is the rate of the i 'th cactus). Competition is modeled as uncertainty, i.e., a cactus i is emptied at each round with some “depleting probability” $0 < s_i < 1$. Furthermore, distances between cacti are

neglected, and in each round, a new cactus can be visited. We restrict attention to purely stochastic strategies that are characterized by a probability vector $p = (p_1, \dots, p_n)$, determining the probability p_i that the bat visits cactus i in each round. We prove analytically that, when cacti are visited often by other bats (modeled by higher depleting probabilities) a purely stochastic strategy based on exploiting a relatively small core will be almost as good as the best strategy. In our simulations, smaller cores can be achieved by increasing α . Indeed, the reinforcement learning literature shows that α can be tuned by an animal, based on the environment [30].

Another interesting result of our analytical analysis (mentioned above) is that the cacti bats should choose for the core are those that have higher refill rates r_i relative to low depletion probabilities s_i . In the analytical model, those cacti with higher r_i/s_i ratios are expected to accumulate larger amounts of nectar. This nicely corresponds to the reinforcement learning algorithm in our simulations, which assigns a large weight to a cactus that contains much nectar upon our bat's visit.

Conclusions

Foraging is a fundamental task that entails many decisions. In this study, we used comprehensive continuous tracking of lesser long-nosed bats in parallel to estimating their available resource to study their foraging decisions as they explore and exploit food resources in their natural environment. Interestingly, a single-parameter model based on reinforcement learning was able to closely reproduce the bats' behavior.

This model has three main characteristics: (1) the emergence of individuals foraging cores through reinforcement learning and competition; (2) exploitation of the cores through stochastic sampling that is dependent on the cacti quality and distances; and (3) exploration of the field in search of new profitable cacti and update of the core accordingly, in parallel to exploitation.

Another foraging strategy that is adopted by many species, including insects and birds, is trapline [31]. In our system, however, we do not see evidence for trapline. In principle, the bats could have used a reinforcement learning component to identify a core with high-quality cacti and then exploit the core using a trapline. However, as we show, it turns out that this combined strategy is inferior to the reinforcement learning strategy (Figure 3G). Perhaps the reason for this is that distances between cacti in a core are small and their qualities still vary. Hence, a strategy that specializes in minimizing traversal lengths, such as trapline, would be less effective than the reinforcement learning strategy that effectively incorporates both qualities and distances despite being memoryless.

Due to the simplicity and effectiveness of the reinforcement-based strategy, we hypothesize that other species facing similar foraging tasks also adopt a similar strategy, a key characteristic of which would be dividing the resources between individuals without direct communication and without a need for aggressive territoriality.

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- KEY RESOURCES TABLE
- RESOURCE AVAILABILITY
 - Lead Contact
 - Materials Availability
 - Data and Code Availability
- EXPERIMENTAL MODEL AND SUBJECT DETAILS
- METHOD DETAILS
 - Tracking bats' movement
 - Movement analysis
 - Foraging behavior analysis
 - Cacti locations and number of open flowers
 - Video observations
 - Nectar refill rate and nectar consumption
 - Estimating the bats' energy consumption
 - The foraging model
 - Alternative models
 - Estimating the overlap of cores
- QUANTIFICATION AND STATISTICAL ANALYSIS

SUPPLEMENTAL INFORMATION

Supplemental Information can be found online at <https://doi.org/10.1016/j.cub.2020.07.079>.

ACKNOWLEDGMENTS

We thank all members of El Pinacate and Gran Desierto de Altar Biosphere Reserve for their generous hospitality, guidance, and help. We thank Rodrigo Medellín's lab members for logistics coordination and Tamir Reuveni, Ruty Goldshtein, Denise Chazan, Isabelle Kains, Tomer Halevi, and Duaa Anabosi for helping process the data. We thank Jonathan Chazan for field assistance, Israel Goldshtein for engineering support, and Bosmat Issahary for graphic assistance. We thank Ofer Feinerman, Ehud Fonio, Inon Scharf, and Eran Amihai for reading and commenting on the manuscript. This research was partially supported by Lev-Zion Scholarship for Outstanding Ph.D. Students from Peripheral Areas, The Minerva Center for Movement Ecology, the Mexican Agency for International Development Cooperation (AMEXCID), and the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (grant agreement no. 648032). The work of Y.E. was supported in part by an Israeli Science Foundation grant number 1016/17.

AUTHOR CONTRIBUTIONS

A.G., Y.Y., and R.A.M. designed the experiments; A.G., M.H., O.E., A.B., T.S., S.C., S.G., and Y.Y. conducted the experiment; and A.G. conducted the analysis and carried out drone imaging and analysis. A.K., S.C., Y.Y., and A.G. designed the model, and A.G. implemented the model. A.K. and Y.E. conducted the analytic analysis of the simplified model. A.G., Y.Y., and A.K. wrote the manuscript.

DECLARATION OF INTERESTS

The authors declare no competing interests.

Received: March 4, 2020
Revised: May 9, 2020
Accepted: July 27, 2020
Published: August 20, 2020

REFERENCES

1. Howell, D.J. (1974). Acoustic behavior and feeding in glossophagine bats. *J. Mammal.* 55, 293–308.

2. Gerardo Herrera, M.L., and Martínez Del Río, C. (1998). Pollen digestion by New World bats: effects of processing time and feeding habits. *Ecology* **79**, 2828–2838.
3. Arends, A., Bonaccorso, F.J., and Genoud, M. (1995). Basal rates of metabolism of nectarivorous bats (Phyllostomidae) from a semiarid thorn forest in Venezuela. *J. Mammal.* **76**, 947–956.
4. Sahley, C.T., Horner, M.A., and Fleming, T.H. (1993). Flight speeds and mechanical power outputs of the nectar-feeding bat, *Leptonycteris curasoae* (Phyllostomidae: Glossophaginae). *J. Mammal.* **74**, 594–600.
5. Medellín, R.A., Rivero, M., Ibarra, A., de la Torre, J.A., Gonzalez-Terrazas, T.P., Torres-Knoop, L., and Tschapka, M. (2018). Follow me: foraging distances of *Leptonycteris yerbabuena* (Chiroptera: Phyllostomidae) in Sonora determined by fluorescent powder. *J. Mammal.* **99**, 306–311.
6. Zimmerman, M. (1979). Optimal foraging: a case for random movement. *Oecologia* **43**, 261–267.
7. Giuggioli, L., Potts, J.R., and Harris, S. (2011). Animal interactions and the emergence of territoriality. *PLoS Comput. Biol.* **7**, e1002008.
8. Stiles, F.G. (1971). Time, energy, and territoriality of the Anna hummingbird (*Calypte anna*). *Science* **173**, 818–821.
9. Wolf, L.L. (1970). The impact of seasonal flowering on the biology of some tropical hummingbirds. *Condor* **72**, 1–14.
10. Feinsinger, P. (1976). Organization of a tropical guild of nectarivorous birds. *Ecol. Monogr.* **46**, 257–291.
11. Linhart, Y.B. (1973). Ecological and behavioral determinants of pollen dispersal in hummingbird-pollinated Heliconia. *Am. Nat.* **107**, 511–523.
12. Kadmon, R., and Shmida, A. (1992). Departure rules used by bees foraging for nectar: a field test. *Evol. Ecol.* **6**, 142–151.
13. Gill, F.B. (1988). Trapline foraging by hermit hummingbirds: competition for an undefended, renewable resource. *Ecology* **69**, 1933–1942.
14. Lihoreau, M., Raine, N.E., Reynolds, A.M., Stelzer, R.J., Lim, K.S., Smith, A.D., Osborne, J.L., and Chittka, L. (2012). Radar tracking and motion-sensitive cameras on flowers reveal the development of pollinator multi-destination routes over large spatial scales. *PLoS Biol.* **10**, e1001392.
15. Janzen, D.H. (1971). Euglossine bees as long-distance pollinators of tropical plants. *Science* **171**, 203–205.
16. Gilbert, L.E. (1975). Ecological consequences of a coevolved mutualism between butterflies and plants. In *Coevolution of Animals and Plants: Symposium V, First International Congress of Systematic and Evolutionary Biology, 1973*, L.E. Gilbert, and P.H. Raven, eds. (University of Texas), pp. 210–240.
17. Janson, C.H., Terborgh, J., and Emmons, L.H. (1981). Non-flying mammals as pollinating agents in the Amazonian forest. *Biotropica* **13**, 1–6.
18. Garber, P.A. (1988). Foraging decisions during nectar feeding by tamarin monkeys (*Saguinus mystax* and *Saguinus fuscicollis*, Callitrichidae, Primates) in Amazonian Peru. *Biotropica* **20**, 100–106.
19. Janson, C.H. (1998). Experimental evidence for spatial memory in foraging wild capuchin monkeys, *Cebus apella*. *Anim. Behav.* **55**, 1229–1243.
20. Watts, D.P. (1998). Long-term habitat use by mountain gorillas (*Gorilla gorilla beringei*). 2. Reuse of foraging areas in relation to resource abundance, quality, and depletion. *Int. J. Primatol.* **19**, 681–702.
21. Lihoreau, M., Chittka, L., Raine, N.E., and Kudo, G. (2011). Trade-off between travel distance and prioritization of high-reward sites in traplining bumblebees. *Funct. Ecol.* **25**, 1284–1292.
22. Thomson, J.D., Slatkin, M., and Thomson, B.A. (1997). Trapline foraging by bumble bees: II. Definition and detection from sequence data. *Behav. Ecol.* **8**, 199–210.
23. Schmidt, J.O., and Buchmann, S.L. (1986). Floral biology of the saguaro (*Cereus giganteus*): I. Pollen harvest by *Apis mellifera*. *Oecologia* **69**, 491–498.
24. Winter, Y., and Stich, K.P. (2005). Foraging in a complex naturalistic environment: capacity of spatial working memory in flower bats. *J. Exp. Biol.* **208**, 539–548.
25. Niv, Y., Joel, D., Meilijson, I., and Ruppin, E. (2002). Evolution of reinforcement learning in uncertain environments: a simple explanation for complex foraging behaviors. *Adapt. Behav.* **10**, 5–24.
26. Börger, L., Dalziel, B.D., and Fryxell, J.M. (2008). Are there general mechanisms of animal home range behaviour? A review and prospects for future research. *Ecol. Lett.* **11**, 637–650.
27. Cvikel, N., Levin, E., Hurme, E., Borisov, I., Boonman, A., Amichai, E., and Yovel, Y. (2015). On-board recordings reveal no jamming avoidance in wild bats. *Proc. Biol. Sci.* **282**, 20142274.
28. Danilovich, S., Krishnan, A., Lee, W.J., Borisov, I., Eitan, O., Kosa, G., Moss, C.F., and Yovel, Y. (2015). Bats regulate biosonar based on the availability of visual information. *Curr. Biol.* **25**, R1124–R1125.
29. Korman, A., Emek, Y., Collet, S., Goldshtein, A., and Yovel, Y. (2020). Exploitation of multiple replenishing resources with uncertainty. *arXiv:2007.09640v1*. <http://arxiv.org/abs/2007.09640>.
30. Behrens, T.E.J., Woolrich, M.W., Walton, M.E., and Rushworth, M.F.S. (2007). Learning the value of information in an uncertain world. *Nat. Neurosci.* **10**, 1214–1221.
31. Ohashi, K., and Thomson, J.D. (2009). Trapline foraging by pollinators: its ontogeny, economics and possible consequences for plants. *Ann. Bot.* **103**, 1365–1378.
32. Tellez, G., Medellín, R.A., Mora, C., and McCracken, G. (2000). Evidence of migration of *Leptonycteris curasoae* in the Mexican tropics. *Bat Res. News* **41**, 143.
33. Egert-Berg, K., Hurme, E.R., Greif, S., Goldstein, A., Harten, L., Herrera M, L.G., Flores-Martínez, J.J., Valdés, A.T., Johnston, D.S., Eitan, O., et al. (2018). Resource ephemerality drives social foraging in bats. *Curr. Biol.* **28**, 3667–3673.e5.
34. Hen, I., Sakov, A., Kafkafi, N., Golani, I., and Benjamini, Y. (2004). The dynamics of spatial behavior: how can robust smoothing techniques help? *J. Neurosci. Methods* **133**, 161–172.
35. Benhamou, S. (2004). How to reliably estimate the tortuosity of an animal's path: straightness, sinuosity, or fractal dimension? *J. Theor. Biol.* **229**, 209–220.
36. Horner, M.A., Fleming, T.H., and Sahley, C.T. (1998). Foraging behaviour and energetics of a nectar-feeding bat, *Leptonycteris curasoae* (Chiroptera: Phyllostomidae). *J. Zool.* **244**, 575–586.
37. Kunz, T.H., and Nagy, K.A. (1988). *Methods of Energy Budget Analysis* (Johns Hopkins University).
38. Boonman, A., Bar-On, Y., Cvikel, N., and Yovel, Y. (2013). It's not black or white on the range of vision and echolocation in echolocating bats. *Front. Physiol.* **4**, 248.
39. Frenzel, J.F. (1993). Genetic algorithms. *IEEE Potentials* **12**, 21–24.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Deposited Data		
Analyzed data	This paper	https://doi.org/10.17632/jxghjsttw.1
Experimental Models: Organisms/Strains		
Lesser long-nosed bat, <i>Leptonycteris yerbabuenae</i>	El Pinacate y Gran Desierto de Altar Biosphere Reserve, Sonoran Desert, Mexico	Taxonomy ID: 700936
Saguaro cactus, <i>Carnegiea gigantea</i>	El Pinacate y Gran Desierto de Altar Biosphere Reserve, Sonoran Desert, Mexico	Taxonomy ID: 171969
Software and Algorithms		
MATLAB R2018a	MathWorks	https://www.mathworks.com/downloads/ ;
R 3.6.1	The R Foundation for Statistical Computing	https://cran.r-project.org/mirrors.html
Pix4Dmapper Pro version 2.2.25	Pix4D	https://www.pix4d.com/product/pix4dmapper-photogrammetry-software
Other		
GPS - Robin	lucid, Israel	N/A
DJI Phantom 3 advanced drone	DJI Science and Technology, China	https://www.dji.com/phantom-3-adv
Full Spectrum POV Cam	GhostStop, USA	https://www.ghoststop.com/Full-Spectrum-POV-Camcorder-p/camcorder-fspovcam.htm
Panasonic Action Cam HX-A1	Panasonic Corporation, Japan	https://shop.panasonic.com/support-only/HX-A1.html
Song meters SM4BAT, Wildlife Acoustics, Massachusetts, USA	Wildlife Acoustics, USA	https://www.wildlifeacoustics.com/products/song-meter-sm4bat

RESOURCE AVAILABILITY

Lead Contact

Further information and requests for resources should be directed to and will be fulfilled by the Lead Contact, Yossi Yovel (yossiyovel@gmail.com).

Materials Availability

This study did not generate new unique reagents.

Data and Code Availability

The datasets and code generated during this study are available at Mendeley, <https://doi.org/10.17632/jxghjsttw.1>

EXPERIMENTAL MODEL AND SUBJECT DETAILS

All experiments were conducted under permit number 443 # 04019/15, 03946/15 14509/16 from the Dirección General de Vida Silvestre (= Wildlife Service). All experiments on *L. yerbabuenae* and *C. gigantea* has been conducted in El Pinacate Biosphere Reserve and in the surrounding Saguaro fields (in the Sonoran Desert in northwest Mexico).

METHOD DETAILS

Tracking bats' movement

During May-June 2015-2017, we successfully monitored the movement of a total of 17 lactating female lesser long-nosed bats (*Leptonycteris yerbabuenae*). We used miniature GPS data-loggers (Lucid, Israel) combined with a synchronized ultrasonic microphone (FG-23329, Knowles) [27] to tag bats roosting in a maternity colony with more than 50,000 lactating bats in El Pinacate Mexico

[32]. A telemetry unit (either LB-2X 0.3 g, Holohil Systems Carp, Ontario, Canada or Pico pip Ag379 0.37 g, Biotrack, Dorset, UK) was attached to the GPS, to enable its recovery after falling off the bat. Tracking bats' movement for up to eight days in a row revealed fundamental aspects of their foraging strategy.

The device's total weight was 4.1 ± 0.1 g (mean \pm SD), which accounted for less than 14% of the average body mass of the bats (the average mass of the bats was: 30.0 ± 2.7 g, $n = 17$). We have already mounted these tags on similar sized bats (e.g., *Rhinopoma microphyllum*, *Myotis vivesi* and *Myotis myotis*) for which we performed several control experiments to test the effect of the extra loading on their well-being and behavior (see detailed information in Cvikel et al. [27] and Egert-Berg et al. [33]). Here too, we made sure that this additional weight did not affect foraging: all tagged bats left the cave and flew to remote foraging sites, sometimes as far as ~ 100 km from their roost where they exhibited vast foraging. We also conducted several control experiments ensuring that flight and foraging abilities of the bats were maintained. To test bats' ability to forage, we captured three non-lactating females, mounted the same GPS tags on their backs, and recorded their behavior while feeding on flowers on a cut Saguaro branch placed inside a flight tent. All bats were able to approach the flowers, hover in front of them and drink (Video S2). Bats were allowed to feed *ad-lib* and were released in good condition at the cave at the end of the experiment. To test the effect of the additional weight on the overall flight duration, we tagged 10 bats with lightweight (1g) telemetry unit (without GPS). The bats that were tagged with heavier GPS devices spent a similar amount of time outside the cave (6.3 ± 1.4 hours, $n = 9$) as the bats that were tagged with light tags (5.9 ± 1.3 hours, $n = 10$, Wilcoxon rank sum test, $p = 0.84$). Additionally, one bat that we tagged and was caught a week later (at the same time of the night), was found in a very good condition and gained 1.3g during this period when it carried the extra tag-weight.

Movement analysis

GPS positions were sampled every 2, 5, 15 or 30 s for different bats (Table S1 summarizes which bats were used in which of the analyses). GPS positioning standard deviation was estimated to be 8 m in the horizontal plane and 11 m in height [27]. Bats' ground speed was estimated as the derivative of distance over time between GPS points that were 15 s apart, and outliers with ground speed above 30 m/s were removed. Bats' airspeed was calculated by subtracting the estimated wind speed in the direction of the flight from the bat's ground speed. Wind data was collected from a weather station located between the cave and the foraging sites (31.679985, -113.304831). The bats' height above the surface was estimated by subtracting the geoid and the elevation above the surface from the altitude above the ellipsoid (provided by the GPS). The geoid height was estimated to an accuracy of 0.001 m using the EGM2008 Geopotential Model and surface elevation was extracted using Google Maps Elevation API. The height over-time was smoothed using the 'LOESS' local regression smoothing filter [34]. Bats' movement was divided to commute and foraging using the straightness index that measures the ratio between the actual path length and the direct distance between two points [35]. The straightness index was calculated for segments of five minutes and revealed a bimodal distribution with a trough at 0.4 separating the two modes. hence, GPS positions with a straightness index lower than 0.4 were considered as foraging and GPS positions with straightness index greater than 0.4 were considered as commute (Figures S2K and S2L). Short flights inside the foraging sites were considered as foraging even if their straightness index was higher than the threshold (0.4).

Foraging behavior analysis

Some of the bats briefly visited secondary foraging sites on the way to the main one (spending no more than 10 minutes in these sites). We restrict our analyses to foraging in the main site.

A visit

A visit to a cactus was defined as flying at a distance of less than 15 m from a cactus, in an altitude of less than 20 m. The video analysis revealed that bats typically lingered in the area of a cactus for a few seconds before visiting it which made GPS-based visit recognition easier. We also validated our visit detection method using the video recordings (Video S1; Figure S2M). Notably, all of our analyses are based on the proportions of visits and not the absolute number of visits, so even if our method for detecting visits under- or over-estimated their number, this should not have affected our results.

Consecutive visits to the same cactus less than 60 s apart were considered as one visit since the video recordings (below) showed that bats tend to fly around the cacti and approach it repeatedly during a visit. In cases where a bat visited a patch with several cacti less than 15 m apart, one of the cacti in the patch was chosen randomly and considered as the visited cactus. For every visited cactus, number of visits, visits duration, distance and time past between consecutive visits were calculated.

We examined the influence of the lower GPS sampling rates on the results, by reducing the sample rate of bats that were sampled every two seconds to five, 15 and 30 s. We found no difference between the bats' behavior using sample rates of every two and five seconds, and corrected the results for only three bats that were sampled at a lower rate (one bat that was sampled every 30 s and two bats that were sampled every 15 s).

The definition of core cacti

Cacti that were visited more than five times during the night by a specific bat were defined as the bat's core cacti. We used a criterion of five revisits to the same cactus because the distribution of revisits drops rapidly above five visits (Figures S1G and S1H).

Cacti locations and number of open flowers

Several approaches were used to assess cactus and flower distribution: (1) cacti locations in specific areas where the tagged bats foraged were identified using (a) drone-imaging (using a DJI Phantom 3 advanced drone with its built-in camera 12.4M pixel, DJI Science and Technology, Guangdong, China); (b) Google Earth satellite and aerial images (in cases where flying a drone in the area was

prohibited). In total, we managed to map the cacti in the foraging sites of eight bats. (2) The flowering of 414 cacti from three cacti fields around El Pinacate biosphere reserve were characterized: 327 cacti (121 flowering cacti) in an area of $\sim 84 \text{ m}^2$ were monitored over 8.3 ± 2.8 days per cactus (22/5/2016–7/6/2016) by remote-imaging using the same drone and camera as above. The drone was flown above the field early in the morning when the flowers are still open. A 3D model of the field was created using Pix4Dmapper Pro version 2.2.25 allowing to count the number of open flowers. Additional cacti were monitored every morning in a manual census during 10.7 ± 3.4 days per cactus (26/5/2015–11/6/2015 $n = 48$ cacti, 6/5/2016–28/5/2016 $n = 23$ cacti, 25/4/2017–2/5/2017 $n = 6$ cacti, 13/5/2017–27/5/2017 $n = 10$ cacti).

Video observations

33 saguaro cacti were videoed for an average of 1.8 ± 1.2 nights ($n = 62$ nights in total) during the flowering season (11/5–1/6/2016, 12/5–18/6/2017) using IR video cameras (Full Spectrum POV Cam, GhostStop, Florida, USA or Panasonic Action Cam HX-A1, H.264 1920x1080, 30fps). Each camera was placed on a two m pole at a distance of two m from the cactus for a period of nine hours (starting at sunset). Bats' echolocation and social calls were recorded in 45% of these observations using Ultrasonic recorders (Song meters SM4BAT, Wildlife Acoustics, Massachusetts, USA) that were placed < 50 cm below the flowers. In cases where the cactus had more than one branch with flowers, we used a separate camera for each branch, so that all approaching bats to all of the cactus flowers could be observed. The videos were fully analyzed and all visits of all pollinators (bats, moths and birds) were denoted. Social interactions between conspecifics were also noted (based on the video or based on the detection of social calls in the audio recordings).

Nectar refill rate and nectar consumption

To estimate the nectar refilling rate throughout the night, the upper part of eight cacti were covered with mesh (20D nylon tulle) to prevent access to their flowers. During two consecutive nights (20–21/5/2017), nectar was extracted from the flowers every 90 minutes (from 21:00 at least until 03:00, during the bats' presence in the field) using a 1 mL syringe, and the amount of nectar was measured ($n = 17$ flowers, 1.4 ± 0.8 flowers per cactus per night). Sugar concentration was measured using Bellingham-Stanley-low-volume Eclipse refractometers (0–50 Brix). We found little variation in sugar concentration ($24.5 \pm 2.7\%$), in nectar rate and in the accumulative nectar quantity (0.73 ± 0.13 ml) at the end of the night, and no correlation between the accumulative nectar quantity at the end of the night and the number of open flowers (Pearson correlation test: $r = 0.1$, $p = 0.6$). The amount of nectar in a flower (ml) can thus be predicted according to the time passed since 21:00 (~ 1.5 hours post-sunset) using a linear equation: $0.13 + 0.102 \times \text{time (hours)}$, $R^2 = 0.89$, $p < 0.001$, $n = 40$ flowers (Figure S1A). To assure that repeated nectar depletion does not affect nectar secretion, we conducted another experiment where we measured the nectar quantity throughout the night in different flowers (each flower was depleted and measured only once). When comparing the two experiments, with and without repeated depletion, we found that nectar secretion rate throughout the night is not affected by repeated nectar depletion (Generalized linear mixed-effects model: $p = 0.381$, $n = 202$ flowers that were measured once, 20 flowers that were repeatedly depleted throughout the night, and depletion status as fixed effect).

To examine how much nectar a bat drinks in a visit, five bats were released in a flight tent for four nights (21:00–03:00), and video-recorded while feeding *ad libitum* on 25% sugar water from 1 mL tubes which were placed inside 10 Saguaro flowers that were mounted on the top of two 20 l water bottles (mimicking the cylinder structure of a Saguaro cactus). After every feeding event, the amount of nectar was measured, and the tube was refilled. Drinking duration of each drinking event was measured from the tube. We found that drinking amount per feeding event could be predicted by the feeding duration, according to a linear equation: nectar feeding amount (ml) = $0.052 + 0.44 \times \text{feeding duration (s)}$ $R^2 = 0.36$, $p < 0.001$, $n = 187$, Figures S2N and S2O). Based on this equation, we concluded that a bat in the field drinks 0.04 ± 0.04 mL (mean \pm SD, $n = 2492$ observations) in an average visit.

Estimating the bats' energy consumption

In order to examine whether the nectar consumed by our bats is enough to cover their energetic demands, we used literature estimates of their energetics. According to Horner et al. [36], a 23 g non-reproductive lesser long-nosed bat spends 1.05, 0.97, 0.32 Watts during commute, foraging and day roosting, respectively. Converting these values to kJ according to the average duration our bats spent in each activity (commuting for 3.1 ± 1.2 hours, foraging for 2.2 ± 1.0 hours and day roosting for 18.7 ± 1.5 hours, $n = 11$) results in a mean energy demand of 40.9 kJ per day. Kunz and Nagy [37] estimated the energy consumption of lactating bats as 50%–100% higher than non-reproductive individuals, hence we estimate the energy consumption of a lactating lesser long-nosed bats as 61.4–81.8 kJ. Since 0.1 mL of Saguaro nectar provides 0.43 kJ [36], lactating bats need to drink between ~ 14 –19 mL of nectar and non-reproductive bats need to drink ~ 9.5 mL of nectar per night.

Based on our GPS monitoring and video analysis, the lactating bats conducted ~ 450 feeding visits on average throughout the night, drinking 0.04 mL on an average visit (according to our video observation and nectar model, Figure S1N). That accumulates to drinking 18 mL of nectar during a full night, nicely corresponding to the calculation above.

The foraging model

In order to unravel the bats' visitation strategy, we developed a mathematical model that simulates a single night of foraging under competition in a space with multiple replenishing food-sources. At each step, each agent (bat) in our model decides which cactus to visit next based on its previous experience and current location. To capture the fact that bats visited profitable cacti more often

(Figure S1M), we modeled this decision process using a rule inspired by reinforcement learning [25], in which the agent learns to prefer rewarded stimuli.

In the simulation, each bat associates each cactus with two parameters: a ‘nectar-weight’ and a ‘score’. The nectar-weight of a cactus depends on the cactus’ nectar yield as experienced by the bat in previous visits, and the score of a cactus depends on both the nectar-weight and the distance of the cactus from the current location of the bat. After every visit to a cactus, the bat updates its nectar-weight (see below) and it then visits the next cactus with probability that is proportional to its score. If higher nectar is found in a cactus, it will have more nectar-weight and hence a higher score, which means that the probability of returning to it can only increase.

In the simulations, the bats start without prior knowledge on cacti quality. At any point in time, bats remember the nectar-weights and positions of only 10 cacti (we chose this number because this is the number of cacti we observed in the core in time slots of 30 minutes (Figure S1J)). Moreover, nectar feeding bats have been shown to have an immense spatial working memory [24]. “In the eyes” of a bat, the nectar-weights of all other cacti are set to 1. Cacti scores are proportional to their nectar-weights and inversely proportional to their distance, that is, the score $S_b(c, t)$ that a bat ‘b’ associates with cactus ‘c’ at time ‘t’ is:

$$S_b(c, t) = \frac{W_b(c, t)}{T_b(c, t) + \tau}, \quad (\text{Equation 1})$$

where $W_b(c, t)$ is the nectar-weight that bat ‘b’ gives to cactus ‘c’ at time ‘t’ (see below), $T_b(c, t)$ is the duration of flight to cactus ‘c’ from the current location of bat ‘b’ (which is proportional to the corresponding distance) and τ is the time duration of a visit.

Similarly to the reinforcement learning framework [25], bats update nectar-weights at each visit based on how much nectar they drank (relative to an average expectation). At time ‘t’, the nectar-weight of the last visited cactus ‘c’ is updated:

$$W_b(c, t) = W_b(c, t - 1) + [N - \text{Avg}] \times \alpha, \quad (\text{Equation 2})$$

where $W_b(c, t)$ is the nectar-weight of cactus ‘c’ at time ‘t’ in the eyes of bat ‘b’, N is the volume of nectar consumed during the visit, Avg is the average consumed volume (set to 0.04ml based on field observations) and α is a scaling parameter equivalent to the learning rate in reinforcement learning.

The list of remembered cacti is updated as follows: if the nectar-weight of a visited cactus exceeds the nectar-weight of at least one of the 10 remembered cacti, it will replace the remembered cactus with the lowest score. In turn, the latter cactus will go out of the list and its nectar-weight will be set to 1. In order to keep exploring, the bats can also visit unremembered cacti (i.e., with nectar weights = 1) but since they do not remember the location of these cacti, they will only visit unremembered cacti within their visual sensing range [28, 38].

The fixed parameters of the model, namely, the numbers of bats and cacti (45 bats with a ratio of 1 bat per 12 flowering cacti), cacti locations, flower distribution and nectar filling rate, were all set based on our field measurements (see below).

The model thus has one free parameter, α - the learning rate. We used an evolutionary algorithm [39] in order to predict the value of α (an evolutionary algorithm searches for a solution that is at equilibrium under the modeled conditions). The value of α in our evolutionary simulations converged to an average of 4.6, with all individuals converging to a range around it (4.6 ± 4.8 Mean \pm SD). Importantly, although we did not fit the data to the model (α was determined evolutionary) the simulated bats with this α behaved very similarly to the real bats (compare the red and black lines in Figures 2B–2F).

The Fix factors of the model

A bat in our simulation arrives at a field with a fixed number of cacti that are distributed over a plane according to an actual field whose cacti-positions have been reconstructed. The simulated field has 1125 cacti and 45 bats. This agrees with our observed bats-cacti ratio estimations: $\sim 200,000$ bats are estimated to roost in the El Pinacate cave using annual bats counting and $\sim 5,700,000$ cacti were estimated to be in a radius of 100 km from the cave based on satellite imaging analysis, see sensitivity analysis for this parameter in Figures S3A–S3F). Each cactus has a certain number of open flowers that are sampled from a real distribution (Figure S1B) and it refills all of them in a constant rate according to our field measurements (Figure S1A).

Initial nectar weight

Bats start at random positions in the field (real bats can fly from one side of the field to the other in ~ 40 s), without prior knowledge on cacti quality (all cacti receive initial equal nectar-weights = 1). We chose the value of 1, since if we choose any other positive constant instead of 1, say c, the system will behave the same by replacing alpha with c times alpha. Since we are aiming to evolve and learn alpha, we will get the same result, only named differently. At any point in time the bats remember the values (i.e., nectar-weights) and positions of only 10 cacti and can visit these remembered cacti and other cacti within their vision detection range. Because a bat always sees additional cacti in its range it will never hold in its memory a cactus with weight smaller than the default value (1), therefore we don’t have to worry about negative weights.

Remembered cacti

The number 10 was chosen because on average, the bats visited 10 core cacti at every moment throughout the night. We divided the night into time period of 30 min and during this period real and simulated bats visited 9.6 ± 0.5 ($n = 8$) and 10.2 ± 1.1 ($n = 45$) core cacti, respectively. Figure S1J). We also show that there is no advantage in remember more than 10 cacti, since even bats that remembering all cacti in the field perform very similar to bats that remember only 10 cacti (Figures S2C–S2F).

Bats' detection range

In the simulation, we assumed that the bats' detection range of a cactus is 50 m. There is very little behavioral data on bats' vision. In a previous study we found that the night vision of *Rousettus aegyptiacus* is similar to humans' [38]. Lesser long-nosed bats have smaller eyes (ca. 0.23 diameter of those of *Rousettus aegyptiacus*) and should thus be less sensitive. To estimate their visual detection range of a cactus, we first estimated humans' detection range and we used this as an over estimation for the bats' detection range. Four humans (adjusted to scotopic vision) were able to detect cacti ($n = 10$) from a distance of 55 ± 9 m. To estimate this, we had them move toward a cactus from a distance of 100 m and report when they first detected it. On the one hand, as explained above, the bats are less sensitive, but on the other hand, their task should have been easier since they observe the cacti from above (unlike the humans) and can search for white flowers. We thus estimated the bats' detection range to be within a range of 30-70 m. To validate that our results are not sensitive to an error of up to 40% in this parameter, we ran our model with cacti detection range of 30, 50 and 70 m (Figures S3G–S3J).

Maintaining core cacti on consecutive nights

Notably, our model assumes that bats are naive at the beginning of the night, while in reality bats' foraging cores on consecutive nights heavily overlap (on a certain night, bats visited $\sim 97\%$ of the core cacti that they visited on previous nights, Figure S3K). This suggests that bats incorporate information gathered on previous nights when forming their foraging core on a given night. Relying on information over consecutive nights probably reduces the costs of finding and mapping the positions of flowering cacti. We also observed that bats reused the same flying pathway between cacti on consecutive nights (Figure S3L), a strategy which could ease navigation in the field. One of the bats that we tracked switched the field it visited while we were tracking it, allowing us to examine its movements in a new field over time. In contrast to all other bats we tracked, in its first nights in the field, this bat spent more time scanning the field and less time foraging, in comparison to the other bats. Over several consecutive nights, it gradually increased the proportion of foraging (Figure S3M). In order to account for this knowledge, we allowed our simulated bats to scan quickly for a short period at the beginning of the night. Thus, at the beginning of the night, our simulated bats explored cacti with a detection range of 250 m. This exploration phase lasts ~ 10 minutes, until each bat acquires 10 remembered cacti. Despite the fact that we neglect this issue of information transfer across nights, our approximated model well captures the general characteristics of bat foraging dynamics.

Learning rate (α)

We used an evolutionary algorithm (using MATLAB's Global Optimization Toolbox) in order to predict the value of α . We ran 30 simulations with 100 generations each and took the average α . We used a population size of 45 bats, with initial values of α that are randomly distributed between 0 and 100. At each generation, 20% of the bats best fitted α (i.e., bats) survived, out of the rest remaining 80%, 62% randomly survived and 18% were crossed with mutation. The values of α always converged within 100 simulations. Nectar consumption was used to estimate fitness.

Modeling aggressive response

A similar model with an aggressive response (aggressive model) provided us the ability to estimate whether social interactions play a role in shaping the foraging behavior of the bats. In this model, simulated bats foraged in the same manner as in the non-aggressive model, except for when two bats simultaneously chose the same cactus. In this case the dominant bat (the bat that has a higher nectar-weight for this cactus) drank from the cactus, and the subordinate bat drinks nothing and is moved near a random cactus 200-500 m from its current location mimicking the result of a territorial agonistic interaction.

Alternative models

We compared the real bats' behavior to two additional models (random and trapline) using the same fixed parameters as in the reinforcement learning model, except for the following changes: In the random model $\alpha = 0$, and thus the bats did not learn from previous experience. In the trapline model, bats used reinforcement learning during ~ 10 minutes at the beginning of each simulation to acquire the first 10 remembered cacti (with $\alpha = 6.1$ that was found using a genetic algorithm). Once the remembered set was established, the simulated bats exploited the core in an ordered manner – always flying to the next closest cactus (as in a trapline). They still occasionally explored additional cacti (according to the same rules of the original model). Finally, we also ran the exact same foraging model with a single bat (and 125 cacti) to compare the nectar consumption of a learning ($\alpha = 4.6$) and a random ($\alpha = 0$) simulated bat.

Estimating the overlap of cores

We used the results of 100 simulations to compute how many bats shared each of the core cacti. For each simulation, we collected the core cacti of all bats into one set S , and then randomly selected a core for each bat, while maintaining the number of core cacti per bat the same as in the specific simulation. That is, if a bat had x cacti in its core in the simulation, we now randomly select x cacti in S to be in its random core. We used a permutation test to compare the overlap in a typical core cactus in these two situations.

QUANTIFICATION AND STATISTICAL ANALYSIS

All statistical analysis conducted using MATLAB R2018a unless stated otherwise. Comparison between groups were conducted using two tailed tests, unless stated otherwise. Non-parametric Wilcoxon rank sum tests were used to examine the influence of the GPS weight on the time bats spent outside the cave because of the small sample size. One-Way ANOVA was used to compare bats' nectar consumption in different foraging models and all other comparisons between models were conducted using Permutation based

t test, to avoid influencing the power of our analysis by changing the sample size (i.e., the number of simulations). We used Linear regression to find the equation of nectar secretion throughout the night and nectar feeding quantity according to feeding duration, and GLM test used to estimate the influence of depletion status (as a fixed factor) on nectar secretion rate throughout the night. Monte-Carlo test conducted using R 3.6.1, ade4 package.