

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Minimal Modeling for Cognitive Ecologists: Measuring Decision-Making Trade-Offs in Ecological Tasks

Permalink

<https://escholarship.org/uc/item/2dn9c12n>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 46(0)

Authors

Forbes, Eden
Beer, Randall

Publication Date

2024

Peer reviewed

Minimal Modeling for Cognitive Ecologists: Measuring Decision-Making Trade-Offs in Ecological Tasks

Eden Forbes (edforb@iu.edu)

Cognitive Science Program, 1001 E. 10th St.
Bloomington, IN 47405 USA

Randall Beer

Cognitive Science Program, 1001 E. 10th St.
Bloomington, IN 47405 USA

Abstract

The complexity of studying behavior and cognitive processes in realistic ecological tasks is a major challenge for cognitive scientists, behavioral ecologists, community ecologists, and the cognitive ecology community that subsumes all these fields. Here we describe a modeling approach that can be used to study the decision-making trade-offs that emerge from the coupling of nervous systems, bodies, and ecological context. To demonstrate the method, we describe an agent that must balance its need to consume resources with its need to avoid predation. We then show how to analyze the resulting behavior through the lens of behavioral trade-off schemas synthesized with neural traces measured during real-time behavior. The employment of model agents will be an important contributor to ecological theory of cognitive processes, and here we hope to convince the reader of that methodological potential.

Keywords: adaptive behavior; cognitive ecology; CTRNNs; dynamical systems; ecological decision-making

Introduction

The last few decades have witnessed many units of study in cognitive science expand from localized brain processes to the brain-body-environment system (Hutchins, 2010) following the development of enactive and embodied philosophies (Varela, Rosch, & Thompson, 1991; Gallagher, 2017), the recognition of adaptive behavior as a critical explanandum to any account of cognition (Chiel & Beer, 1997), and the resulting identification of new potential systems of study beyond humans and “higher” organisms (Hutchins, 1995; Lyon, 2015). As such, more and more work aims to describe the effects of cognitive abilities on adaptive behavior in ecological context. One burgeoning field that tackles these issues is cognitive ecology, which aims to describe and explain the cognitive processes that underlie adaptive behavior in real animals (Dukas & Ratcliffe, 2009). Cognitive ecology appeals directly to the greatest source of examples we have of successful brain-body-environment systems: the flora and fauna of the natural world. In appealing to the behavior of real organisms in response to issues of survival and reproduction, the cognitive ecologist aims to circumvent the pitfalls of defining cognitive phenomena abstractly and then fitting an artificial task to that definition. This is to some degree a practical move, given the regular disagreement in the mind sciences of what counts as “cognitive” in the first place (Allen, 2017). Time is better spent engaging with problems that real organisms encounter and use that as our foundation as a field.

Furthermore, an ecological approach encourages the development of a more robust theory of cognition that transcends the idiosyncracies of a particular species, which is especially important given the historically anthropocentric benchmarks of cognitive science (Van Duijn, Keijzer, & Franken, 2006).

While this tack self-evidently brings to bear real examples of adaptive behavior and cognitive phenomena, those phenomena of interest in the wild demand significant breadth and volume of data including measures of cognitive and behavioral performance, associated control structures, and relevant ecological conditions (Mettke-Hofmann, 2014). Researchers have generally responded to this hefty demand in one of two ways: they either attempt to extract general cognitive archetypes from a synthesis of existing data that works across systems, or they attempt to interrogate one system, often a particular species, relying on accrued data and observation in an attempt to faithfully describe behavior and its mechanism in one taxon. Cognitive archetypes are useful in theory but tend to be loosely defined given that they are deeply circumstantial, sacrificing biological realism to create models with abstracted premises that still provide some precise results. Conversely, case studies of specific species provide more predictive power within that species, but often sacrifice any generality across ecological conditions and almost always fail to extend to even closely related species. As such, it will be remarkably difficult to obtain any general results from, say, a single model organism given the remarkable complexity of the interplay between cognitive processes, physiological processes, and ecological processes. It appears that tangible progress towards more comprehensive theory obviates other approaches that can be validated by experimental results but are not severely limited by time and resource constraints.

Levins (1966) champions a third approach: sacrifice numerical precision for realism and generality. Step away from exact prediction of an observed system and attempt to derive conclusions that operate within realistic biological constraints while yielding useful insight across various systems. This is not to suggest sacrificing the rigor of analysis, but rather any expectation that the results derived from that analysis will quantitatively apply to a specific observed biological system. Two related approaches that have emerged in cognitive science along these lines are the adaptive behavior and evolutionary robotic programs. Beer and Gallagher

5841

(1992) and (Harvey, Paolo, Wood, Quinn, & Tuci, 2005) describe frameworks for generating analytically tractable models of behavior through the evolution of nervous systems that control a specified body in a certain task environment. This approach, in which the mode of behavior is not specified but instead allowed to emerge from an idealized coupled brain-body-environment system, has been fruitful in modeling some of the ethology of simple animals (Izquierdo & Beer, 2013) and "minimal" cognitive processes such as object recognition and relational object discrimination (Beer, 1996; Slocum, Downey, & Beer, 2000; Beer, 2003). Researchers benefit from the transparency of these models; the entire system including the implementation and activity of the nervous system, body, and environment can be directly interrogated to investigate how those parts work together in concert to realize the behavior. Furthermore, comparative questions about possible differences in nervous structure, morphology, or environmental circumstance are all explicitly formalized by modeling choice rather than left to verbal interpretation. This combination of analytical tractability and the multiple scales of biological phenomena that can be accounted for is unique, powerful, and certainly of use to cognitive ecologists given its ability to account for those tripartite relationships that the field is interested in.

Of course, a truly robust understanding of a field or phenomenon requires that a variety of scientific approaches – those that sacrifice precision, those that sacrifice generality, and those that sacrifice realism – converge on similar conclusions. No modeling method should necessarily be prioritized over any other, all three have their place. However, when it comes to cognitive ecology, we believe that the theoretical dimension deserves more attention and rigor, and that one way to provide that attention and rigor is to combine methods from research in adaptive behavior and in theoretical ecology. Here, we hope to back this claim by outlining a methodological pipeline that captures the unit of study that cognitive ecology is interested in: analysis of an ecologically relevant behavior that considers information about the agent's physiology, including its nervous system, and the ecological circumstances. We will exemplify these methods and analyses using a rudimentary example that is not meant to represent any specific biological system, but can be complexified to reflect specific ecological systems in future applications. As a result, we hope to convince the reader of the utility of the method. We will place specific emphasis on how to characterize the behavior of agents in these models as a product of ecological trade-offs that can be identified mathematically in our agents' evolved nervous systems.

Methods

Task Environment In our model, agents are situated in a continuous one-dimensional world that is populated by food sources and a single predator agent. Food sources are initialized at random positions throughout the environment up to a specified carrying capacity K . This carrying capacity varies

across evolutionary stages incrementally between a lower and upper bound (K_{min} ; K_{max}) so agents are required to solve the task in environments of decreasing abundance. In a trial, food sources are depleted when crossed by the agent according to the feeding rate of that agent which then assimilates a certain amount of nutrient according to its feeding efficiency. Food has no regenerative property, so any reduction of a food source is permanent. Once a food source is fully depleted, it is removed from the environment. At each timestep, there is a random chance scaled by an implicit food growth rate and the current difference between food population and carrying capacity for a new food source to appear at a random position in the environment. As such, food sources will slowly rebound to carrying capacity K if depleted.

The single predator is a hard-coded agent initialized at a random position in the environment. It is equipped with a Gaussian sensor for detecting the prey agent:

$$I = -ae^{-\frac{x^2}{2\omega^2}} + ae^{-\frac{y^2}{2\omega^2}} \quad (1)$$

where x and y are given by the distance to the prey agent on the left and right of the predator, a scales the upper limit of the Gaussian, and ω specifies the width of the Gaussian spread. The sensor is inverted at the predator's position and signed such that the leftward distance will yield a negative signal and the rightward distance will yield a positive signal. If this sensor value rises above a certain threshold, the predator is considered to be "tracking" the prey and moves directly according to that sensor readout multiplied by a scalar gain value that determines the predator's maximum speed. Otherwise, the predator moves according to one of three default conditions that vary across trials during evolution, either drifting to the left at half of its maximum gain, drifting to the right at half of its maximum gain, or staying still. If a predator crosses the evolved prey agent during its movement, it consumes the agent and the trial ends. Only one predator is used due to constraints of a one-dimensional environment. In the current setup, multiple predators can easily conduct a pincer maneuver such that they pin the evolved agent between them with no chance of escape. Consequently, the scale of the environment lies more on the order of microhabitat rather than true ecological range in which organisms may encounter many predators. More nuanced predator-prey interactions are left for future work.

Agent Physiology Crafting agents for study in an ecological paradigm requires explicit description of their basic metabolic processes. Metabolism is a lynchpin for creating interesting ecological tasks for several reasons. Focusing on the uptake, allocation, and expulsion of material by biota creates a common currency to discuss individual life history, population interactions, and ecosystem processes across taxonomic and environmental differences (Brown, Gillooly, Allen, Savage, & West, 2004; Loreau, 2010). This argument not only applies in vitro but in silico as well; formalizing metabolic processes allows morphologically, neurally,

or otherwise dissimilar agents to deal with similar fundamental tasks in a shared environment.

Here, agents consume resources that they find in the environment at a certain rate and efficiency through movement coordinated by the evolved nervous system. The assimilated resources enter the metabolic store of the prey agent while that store is constantly drawn upon for realizing the metabolic processes necessary to keep the agent alive and active in its environment. There are two critical thresholds imposed on the metabolic store. If the store falls below a value of 0.0, the agent is presumed to have died and the trial ends. Conversely, if the store rises above a certain birth threshold, the agent is scored for giving birth and contributes a certain amount of its metabolic store to its theoretical child.

Agent Nervous Systems The prey agent’s nervous system is modeled using a combination of pre-specified sensors and a 3-node continuous-time recurrent neural network (CTRNN). Each node of the CTRNN takes on the following standard form as detailed by Beer and Gallagher (1992):

$$\tau \frac{dy_i}{dt} = -y_i \sum_{j=1} w_{ji} \frac{1}{1 + e^{(\theta_j - y_j)}} + I_i \quad (2)$$

Agents move according to the difference between two of the neuron outputs in the CTRNN ($\sigma(y_2 + \theta_2) - \sigma(y_1 + \theta_1)$) scaled by a gain parameter. Three sensors that are fully connected to the CTRNN respond to the agent’s metabolic state, the location of food, and the location of predators. The state sensor is a direct read of the metabolic store as described in the agent’s physiology and in turn can take on the same range of values as that store. The food and predator sensors are determined according to the closest source’s distance according to the same Gaussian sensor used by the predator in Equation 1.

Evolutionary Algorithm While the nervous system described here could likely have been solved analytically for this task, we use an evolutionary approach to show how to work with agents that have been evolved on arbitrarily difficult tasks. Previous work has shown that evolutionary approaches help avoid assumptions from the researcher about what the optimal solution to a task should look like and instead allows us to explore the full hypothesis space of solutions (Beer & Gallagher, 1992; Izquierdo, 2019). As such, a genetic algorithm was employed to search for possible solutions to the specified task environment. The algorithm had access to all parameters of the CTRNN nodes and their connections (τ, θ, w) as well as the weights between the sensors and the CTRNN nodes (w_S). With a 3-N CTRNN and 3 sensors, this totals to 24 free parameters that are adjusted during evolution. Each genotype was tested twice on each of the three predator conditions for a total of six trials per generation. Each trial continues up to a maximum number of timesteps unless the agent is consumed or starves, in which case the trial ends. For a given trial, the prey agent’s fitness was calculated as:




Figure 1: All 100 evolutionary runs conducted. In total, 70 of the 100 evolutions solved the task at each value of K . Here the top 19 successful agent evolutions are shown in color, all of which tied for the highest fitness value of 1.03.

$$Fitness = \frac{Time\ Survived}{Maximum\ Time} + \frac{Number\ of\ Births}{K * 100} \quad (3)$$

So, most of the fitness is determined by how long the agent persists in the simulation, while a small bonus is attributed for potential population growth. Fitness for a given generation is the minimum fitness value across the six trials. Notably, the fitness function selected by the researcher has significant impact on the resulting behavior, and is one of the most sensitive points for including biological realism in the model. We make no claim that this fitness function is meant to represent any particular biological reality, but here was used to generate agents with interesting trade-offs in the task environment.

Lastly, a staged version of evolution was employed such that each time the minimum trial fitness rose above 1.0, implying high likelihood that the agent survived the entire run, the test carrying capacity K was decreased by 1 starting from K_{max} down to K_{min} . So, each agent is exposed to harsher environments throughout evolution with an assumption that adjustment to harsher conditions will still imply success in more relaxed cases. If an agent achieves minimum trial fitness above 1.0 at K_{min} , it is considered a successful agent and is tested at a final carrying capacity K_{rest} for the remaining generations. We increase K at the end of evolution to allow for a slightly broader space of possible solutions all derived from the same harsh (K_{min}) constraints¹.

Analysis

A plot of all 100 evolutionary runs of individual prey agents is given in Figure 1. The best runs, those that solved the task at K_{min} and tied for the highest fitness, are shown in color. While 70 of the 100 evolutionary runs solved the tasks at every value between K_{max} and K_{min} , 19 agents tied for the greatest fitness

¹Code and parameters for both simulation and analysis in this study can be found at <https://github.com/eforbes24/CogSci2024.git>




Figure 2: *Sample behavioral traces of Agent 15 in the task environment. Blue indicates the evolved agent, green indicates food sources, and red indicates the predator agent. At the parameters used for this example, evolved agents have a greater maximum speed than the predator, and as such can find a strategy where they keep the predator close by and accelerate to collect food before the predator can catch up. This strategy is not meant to replicate any particular animal strategy, but will be used here as a case study of the proposed methods.*

value (1.03) at K_{test} . Below we will discuss only those top 19 agents with special emphasis on the analysis of Agent 15.

Quasi-Static Approximation of Behavioral Trade-Offs

While our agents are rudimentary, they are still subject to the same ecological constraints that we would expect a real organism to encounter. In particular, each agent must balance its need to consume resources with its need to avoid predation. Figure 2 shows Agent 15 navigating this trade-off in the task environment. As cognitive scientists, we are interested in how that trade-off is mediated and the relevant decisions are made to keep the agent alive. While in real animals we do not have access to all of the information needed to answer those questions, here we have access to the entire nervous system, body, and environment, and as such can interrogate those questions directly using the tools of dynamical systems theory and computational neuroethology.

One productive approach is to examine how the agent’s sensory input translates to their movement in the environment and how the agent responds to different combinations of those inputs. Again, each agent has access to three channels of sensory information: the location of food, the location of the predator, and their current internal state. We can construct trade-off schemas that describe how the agent responds to changes in each of those senses via a quasi-static approximation (for examples, Beer (1995) and Agmon and Beer (2014)) derived from the agent’s nervous system. Take for example the nervous system of Agent 15, shown in Figure 3A. The ODEs that describe the CTRNNs which govern the nervous dynamics can be solved for a certain sensory situation if we consider those sensory inputs as fixed parameters

of the system. In turn, we can take the difference of the output values of neurons 1 and 2 to determine the equilibrating direction and magnitude of agent movement at that fixed sensory combination.

Figure 3B shows a projection of the resulting trade-off schema for Agent 15 at a certain self-state (1.5), from which we can start to describe the trade-offs that agent uses. All our successful agents take the presence of the predator very seriously; movement choices largely correspond with the sign of the predator signal (along the y-axis). This is perhaps unsurprising; given that contact with the predator immediately ends the run, careful predator avoidance is necessary for success at the task. Food acquisition is somewhat more complicated, but again follows a similar pattern across our best agents. If we focus on circumstances where there is no predator (0 on the y-axis in Figure 3B), oscillation around food sources is possible due to consistent rightward movement when food is close to the right (as the food sensor approaches 2.0) and consistent leftward movement when food is close to the left (as the food sensor approaches -2.0).

Each successful agent varies somewhat in the exact realization of the trade-off schema, but these two features are conserved in all 19 best fit cases. They also work synergistically in a similar manner. All successful agents find the predator and then keep it at a certain distance as they traverse the environment in search of food. This allows them to then accelerate when food is detected and give them a small window of time to consume the food while the predator catches up. In some cases, agent acceleration will cause the predator to lose sight of the agent. However, the prey’s baseline direction of exploration (movement direction at (0,0) in Figure 3B) always moves it back in the predator’s direction, thus resuming the monitoring behavior.

Synthesis of Neural Traces and Trade-Off Schemas

While the strategies of the top agents all were behaviorally similar, there were no mechanisms in the nervous system that were shared across all agents. As such, it’s important to show that we can examine how each respective nervous system realizes the behavior in more detail. In real time, agents motor output will rarely achieve the exact equilibria values described by the trade-off schema; any action the agent takes feeds back into its sensors and changes its position in the schema. That said, the approximation is an informative tool if the agents are integrated at a fine enough interval. Figure 4 juxtaposes the trade-off schema and recorded neural traces of the agent during real-time behavior so we can examine how the trade-off schema is actually employed during a particular behavioral sequence. Here, as the agent approaches a food source (i) it accelerates, increasing its distance from the predator. This trick exploits the Gaussian sensor employed by the predator; the agent remains just in the predator’s view, but at such a distance that the predator moves very slowly towards it while the agent eats (ii) (if the predator “lost sight” of the agent entirely, it would go at its much faster drifting pace toward the agent in some conditions). Once the food is con-




Figure 3: (A) The evolved nervous system of Agent 15. $y_1 - y_3$ are neurons modeled by the CTRNN equation while s_f , s_p , and s_s are the sensors for the food, predator, and self-state respectively. Black arrows indicate excitatory connections while red punches indicate inhibitory connections. (B) The trade-off schema derived from solving the equations that govern the nervous system at various combinations of fixed sensory inputs. This quasi-static approximation allows us a first pass at the behavioral trade-offs that the agent employs in the task environment. Positive (more yellow) values indicate rightward movement, negative (more purple) values indicate leftward movement, and neutral (green) values indicate little to no movement. Each quadrant (Q1-Q4) depicts one possible situation the evolved agent may find itself in, each graphically depicted on the right with blue circles indicating the agent, red triangles indicating the predator, and green posts indicating food sources. So, for example, in quadrant 1 when food is to the agent's left and the predator is to the agent's right, the schema predicts leftward movement unless the sensory inputs are very low.

sumed, the prey agent finds itself lacking any strong sensory input from the predator and resumes motion back toward the predator until it comes into view and their normal traversal through the environment resumes (iii).

The neural traces reveal which areas of the trade-off plots are employed during each of these behavioral phases and which nodes of the CTRNN change in correspondence with change in position on the trade-off schema. This goes to show that much of the trade-off schema may never be visited by successful agents during their actual behavior. As such, blindly relying on the trade-off surfaces alone may lead to erroneous conclusions if one attempts to extrapolate specific behaviors across the entire schema. Examining the schema's relationship to neural traces is also a useful extension if we are interested in identifying features of the nervous system that are similar across successful agents or how circuits differ across various evolutionary conditions. For now, more nuanced analyses of the nervous dynamics are left for future work, but we are left with the impression that this combination of behavioral trade-off schemes and neural traces will be a useful tool for those interested in the relationship between nervous systems and behavior in ecological context.

Discussion

Here we have described the evolution of a nervous system situated in a greater brain-body-environment system that exemplifies ecological decision-making in the navigation of the trade-off between foraging and predation. Furthermore, we

have shown how one may approach visualizing and analyzing not only the resulting nervous system but the actual decision-making schemas that fall out from their coupling with the body and environment. While the example we lean on here is somewhat simple, it is just one worked example of a methodology that can be applied to a variety of other tasks and environments and which may be especially powerful in comparative cases that are often of interest to decision-making researchers and cognitive ecologists. Additionally, quasi-static approximations can be taken from a system of any dimensionality, so this approach can be applied to more or less complex sensory and nervous systems. In any case, theory in cognitive science and cognitive ecology become much more analytically tractable; being able to analyze the trade-offs an agent uses in an ecological task moves us beyond verbal description or hypothesis and into formal theory.

It furthermore is a step forward for the evolutionary robotic approach in that it moves beyond single agents and towards entire communities of agents. For example, Nolfi and Floreano (1998) describe the co-evolution of behavior in a prey and a predator, a version of which could be implemented in this same task environment. Not only would this afford an examination of the co-evolution of decision-making trade-offs, but it would also allow us to examine the ecological consequences of those trade-offs. In many cases, the consequences of certain behavioral strategies on a given ecology remain opaque due to the difficulty of collecting or even systematizing comprehensive behavioral and ecological data.




Figure 4: Application of neural traces to the trade-off schema that show how the agent navigates the foraging/predation trade-off. (A) Trace recordings of the sensors and neurons of the agent with (B) the corresponding real-time behavior divided into (i) food approach, (ii) food consumption, and (iii) predator monitoring. Light gray, medium gray, and dark gray are the outputs of $y_1 - y_3$ respectively, blue is the self-state sensor, green is the food sensor, red is the predator sensor, and orange is the motor output of the agent. (C) Here the sensory traces are superimposed on the trade-off schema, with the traces colored by the actual motor output of the agent. Each panel corresponds with the three phases described above (i-iii). Slight changes in the trade-off schemas result from the changing internal state of the agent before, during, and after feeding. In phase (i), as the agent detects and begins to approach the food source, it pulls away from the predator via the excitation of y_1 from the food sensor and the subsequent loss of excitation in y_1 as the predator sensor decreases. When the agent reaches the food, the predator is at the periphery of its sensory range. The asymmetry in the trade-off scheme underlies the oscillation in phase (ii) where the rapid sign-switching of the food signal pulls the agent back and forth across the food source. In phase (iii), the absence of a food signal leads the agent back to the exact center of the trade-off schema, and soon thereafter it begins traversing back in the direction of its original approach and the predator. Upon reaching the predator, the agent again switches direction and keeps it at the same following distance as before.

Fortunately, working with simulated agents uniquely affords the ability to collect data on the ecological consequences of behavior at every demographic combination the system could exhibit. As such, the system of differential equations that describes an ecology that includes certain evolved behaviors can be directly calculated and analyzed using the tools of dynamical systems theory. We can in turn ask questions that point in both directions; how does ecology impact the evolution of behavior and how does the resulting behavior impact the ecol-

ogy? It is these causes and consequences that are of particular interest to cognitive ecologists, and the ability to examine them comprehensively will be a step forward for that field and for the study of cognition "in the wild" (Hutchins, 1995).

Acknowledgments

The authors would like to thank the reviewers for their helpful comments and suggestions as well as Connor McShaffrey and Gabriel Severino for their helpful discussions.

References

- Agmon, E., & Beer, R. D. (2014). The evolution and analysis of action switching in embodied agents. *Adaptive Behavior*, 22(1), 3–20.
- Allen, C. (2017, November). On (not) defining cognition. *Synthese*, 194(11), 4233–4249.
- Beer, R. D. (1995, March). On the dynamics of small continuous-time recurrent neural networks. *Adaptive Behavior*, 3(4), 469–509.
- Beer, R. D. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior.
- Beer, R. D. (2003). The dynamics of active categorical perception in an evolved model agent. *Adaptive behavior*, 11(4), 209–243.
- Beer, R. D., & Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive behavior*, 1(1), 91–122.
- Brown, J. H., Gillooly, J. F., Allen, A. P., Savage, V. M., & West, G. B. (2004, July). Toward a metabolic theory of ecology. *Ecology*, 85(7), 1771–1789.
- Chiel, H. J., & Beer, R. D. (1997). The brain has a body: adaptive behavior emerges from interactions of nervous system, body and environment. *Trends in neurosciences*, 20(12), 553–557.
- Dukas, R., & Ratcliffe, J. M. (Eds.). (2009). *Cognitive ecology II*. Chicago: University of Chicago Press.
- Gallagher, S. (2017). *Enactivist interventions: Rethinking the mind*. Oxford University Press.
- Harvey, I., Paolo, E. D., Wood, R., Quinn, M., & Tuci, E. (2005, January). Evolutionary robotics: A new scientific tool for studying cognition. *Artificial Life*, 11(1-2), 79–98.
- Hutchins, E. (1995). *Cognition in the wild*. MIT press.
- Hutchins, E. (2010). Cognitive ecology. *Topics in cognitive science*, 2(4), 705–715.
- Izquierdo, E. J. (2019). Role of simulation models in understanding the generation of behavior in *c. elegans*. *Current Opinion in Systems Biology*, 13, 93–101.
- Izquierdo, E. J., & Beer, R. D. (2013). Connecting a connectome to behavior: an ensemble of neuroanatomical models of *c. elegans* klinotaxis. *PLoS computational biology*, 9(2), e1002890.
- Levins, R. (1966). The strategy of model building in population biology. *American scientist*, 54(4), 421–431.
- Loreau, M. (2010). *From populations to ecosystems: Theoretical foundations for a new ecological synthesis*. Princeton: Princeton University Press.
- Lyon, P. (2015). The cognitive cell: Bacterial behavior reconsidered. *Frontiers in microbiology*, 6, 264.
- Metzke-Hofmann, C. (2014, May). Cognitive ecology: Ecological factors, life-styles, and cognition. *WIREs Cognitive Science*, 5(3), 345–360.
- Nolfi, S., & Floreano, D. (1998, October). Coevolving predator and prey robots: Do “arms races” arise in artificial evolution? *Artificial Life*, 4(4), 311–335.
- Slocum, A. C., Downey, D. C., & Beer, R. D. (2000). Further experiments in the evolution of minimally cognitive behavior: From perceiving affordances to selective attention. In *From animals to animats 6: Proceedings of the sixth international conference on simulation of adaptive behavior* (pp. 430–439).
- Van Duijn, M., Keijzer, F., & Franken, D. (2006). Principles of minimal cognition: Casting cognition as sensorimotor coordination. *Adaptive Behavior*, 14(2), 157–170.
- Varela, F. J., Rosch, E., & Thompson, E. (1991). The embodied mind. *The embodied mind: Cognitive science and human experience*.