

Cockroaches, Drunkards, and Climbers: Modeling the Evolution of Simple Movement Strategies Using Digital Organisms

Wesley R. Elsberry, Laura M. Grabowski, Charles Ofria, and Robert T. Pennock
(E-35 Holmes Hall, Michigan State University, East Lansing, MI 48825, USA)

Abstract—Even the simplest of organisms may exhibit low-level intelligent behaviors in their directed movements, such as in foraging. We used the Avida digital evolution research platform to explore the evolution of movement strategies in a model environment with a single local resource that diffuses to produce a gradient, which organisms have the ability to follow. Three common strategies that evolved, *Cockroach*, *Drunkard*, and *Climber*, exhibit how both environmental constraints and historical contingency play a role in the emergence of intelligent behaviors. The evolved programs are also suitable for use in controllers on robots.

I. INTRODUCTION

There are a variety of choices for the scale of problem to examine when it comes to issues of movement of agents. At the smallest scale, there are problems of kinematics, where matching patterns of activation of limbs or wheels to useful agent motion is the focus. At the largest scale, there are problems of characterizing particular patterns of movement by agents, and exactly how an individual agent coordinates its internal states to obtain motility is not relevant. Our research takes an approach at an intermediate level, seeking to elucidate how evolutionary processes can result in individual control of existing movement capabilities in order to intelligently exploit environmental resources.

There have been many approaches in computational systems that touch upon the movement of agents. The question of interest in each case makes a difference as to what capabilities are given to such agents, and what properties or capabilities are altered through the time course of the computation. The question of interest for us lies in how intelligent behavior arises through evolutionary processes, so our approach does not specify any particular way of generating such behaviors, but leaves that open to exploration via those processes.

Our interest requires a more open-ended approach than is often used in studies of computational intelligence. In most cases, there is a specific function of interest to be accomplished, and the means or process by which that function is acquired is of less interest than the fact of either solving, or approximately solving, it. Instead, in looking at the evolution of intelligent behavior, our primary interest is in finding out by what means less capable agents give rise to those able to appropriately exploit prevailing conditions.

Given that a rich repertoire of behaviors in biological organisms concerns *movement*, where examples abound ranging from chemotaxis in bacteria through route planning for human delivery drivers, exploring the use of movement by

evolving agents can open up many research questions that are directly comparable to work within biological systems.

Computational research touching on movement of agents spans many different fields. Movement may not be modeled at all, but simply assigned a cost value, as in work in artificial neural systems applied to the traveling salesman problem [1]. For many purposes, specifying one or more fixed movement strategies, possibly modified stochastically by parameters, suffices. Such approaches have been seen in individual-based models (IBM) in ecology (e.g., Roese *et al.* [2]) and artificial life (e.g., Echo [3] and Gecko [4]).

In evolutionary robotics, interest in the movement of agents also covers a variety of levels. These range across studies looking at environment recognition given an initial set of behaviors [5], evolving kinematics for particular controller anatomies ([6], [7]), evolving neural controllers for movement [8], and direct evolution of a Turing-complete binary code to accomplish an obstacle-avoidance task with a robot [9].

Our hypothesis is that given a means of extracting information from the environment, processing it, and having the capability for movement in the environment, evolutionary processes can discover *effective methods* for exploiting spatially-distributed resources. For the purpose of testing this hypothesis, we use the Avida artificial life software platform. Our approach provides organisms with the capacity to make single-step movements, to alter their heading randomly to another of a limited set of discrete headings, and to retrieve information from the environment about local differences in resource availability, but with no information about how to use these capabilities in order to best make use of environmental conditions.

An important consideration for our approach to this problem is that we be able to elucidate the patterns of emergence by which evolutionary computation finds intelligent behaviors. We can usefully classify movement strategies based upon the acquisition and use of information from the environment, where effective and efficient collection and application of information in decision-making is a key property of intelligence. Our goal is to investigate the patterns of emergence in the evolution of such behavior, taking note of both theoretical and applied concerns. In the simple environment we have modeled, an optimal strategy exists, other things being equal, namely, direct gradient ascent. At the opposite extreme, a random walk is the worst strategy. Not moving provides neither benefit nor cost, but a random walk can actually cost more than the benefit

provided. Other movement strategies may not be optimal but nevertheless may be good enough. What behaviors evolve and what effects do historical contingency and environmental constraints have? Can evolution produce effective methods at a level of abstraction suitable for real-world application with different sensory modalities?

Toward that last concern, we also worked with robots to instantiate various of the evolved programs found in our experiments.

II. METHODS

In a preliminary step for exploring the evolution of intelligent behavior with choice expressed by movement, we extended the proven artificial life platform Avida ([10], [11], [12], [13], [14]). Avida provides some features that are useful for implementing movement and the specification of spatially distributed resources, but other features must be either modified, or added as new to the system. The primary property of each digital organism in Avida, or Avidian, is its *genome*, a circular list of instructions from a Turing-complete instruction set. Avidian instructions are executed on a virtual CPU. The CPU is outfitted with three registers and two stacks. The instruction set permits labels, conditional evaluation, looping, arithmetic, logic operations, and everything needed to permit *self-replication* of the Avidian genome and generating an offspring Avidian thereby. Avida does not use explicit fitness functions, instead permitting the experimenter to associate *merit* with the completion of certain tasks or emission of particular behaviors. Avidians with higher accrued merit are awarded a proportionally greater share of available time-slices, meaning that they are able to process instructions faster and may replicate more quickly than Avidians with lower merit. Time in Avida is measured in *updates*, where each update corresponds to a number of instructions performed, on average, by each Avidian in the population. The methods used in this research overlap with the research of Grabowski *et al.* (2008) [15], where further details on methods can be seen.

We added three instructions to Avida: `tumble`, `move`, and `sense-diff-facing`. Avida uses a discrete two-dimensional grid for placement of Avidians. Each Avidian has a *facing* that orients it toward an adjoining grid cell. The `tumble` instruction changes the current facing to a random facing toward a different adjoining grid cell. The `move` instruction causes the Avidian to move to the grid cell it faces. If another Avidian is in that grid cell, the two swap places. Facings are preserved across moves, or altered if the facing is not legal in the new grid cell. The `sense-diff-facing` instruction compares the amount of a resource in the current grid cell and the grid cell faced, placing that difference in register B.

To model a simple environment with a local resource with diffusion, a resource peak was specified with concentrated resource availability in a local region, and a gradient of lesser resource availability established across the remainder of a bounded grid (see Figure 1). No merit was accrued by any organism without movement, but each step taken

on `move` incurred a 2% movement cost in merit. None of the logic tasks used in previous work with Avida were rewarded. Instead, a new task, *dont-care*, was associated with the spatially-distributed resource, and evaluation only occurred when an Avidian executed a `move` instruction. Avidians accrued more merit when they moved to grid cells with higher levels of the resource. A population cap, or *carrying capacity*, was implemented, such that on the birth of an Avidian, another randomly selected Avidian other than the parent was killed to keep the population within the specified carrying capacity.




Fig. 1. Resource distribution, shown in both 3D and 2D projection. The scale shows the mapping between color and amount of resource per grid cell.

In our first experiment, 100 runs were conducted, each seeded with a basic replicating Avidian with no other capabilities. The grid size was 111x101 grid cells. We used a bounded grid model, and chose the size of the sides to be relatively prime so that there was less chance that our Avidians could exploit some periodicity implicit in the grid size. Runs lasted for 500,000 updates, with each update representing three instructions executed per organism on average. The population size was capped at 200 Avidians, leaving slightly over 98% of grid cells unoccupied at any given update in the run. There was no depletion of the spatially-distributed resource during the run. Information about the Avidian population, the resource distribution, and a tally of *visits* by Avidians to each grid cell was kept from each run, with summaries logged every 5,000 updates.

From among the movement strategies that were observed in these runs, a selection of evolved programs that approximated a gradient ascent method (the class of *Climber* movement strategies) that performed well were used to test the generality of the programs. One aspect of intelligent behavior is that such behaviors provide adaptive responses to novel stimuli. For this test, evolved *Climber* programs were injected into Avida runs with random placement of the resource peak.

The transition probabilities derived from the 100 runs of the first experiment were used to simulate sets of runs of various lengths via Monte Carlo methods. A Perl script was used to read in the transition probability data, then iterate such that beginning from the start state the simulation would select the following state based upon a pseudo-random

number and the relevant transition probabilities. Each iterated step corresponded to 5,000 updates in an Avida run.

We translated evolved programs from Avida to control an iRobot Create robot, using software derived from a previous effort [16]. Changes from that work included a re-organization of code generation to fit larger evolved programs into the limited memory space of a robotic controller and completion of the virtual CPU instruction set for the robot context. The Create Command Module, an Atmel ATMEGA16-based microcontroller system, was used for hosting our program. The program consists of a small robot driver program, the genome data, and a minimal version of the Avida virtual CPU, including its stacks and registers. The Create was outfitted with two CdS cell light sensors, one at the front and the other at the rear of the robot. For most Avida instructions, the only things that change are states of the virtual CPU. For the three added instructions, `move`, `tumble`, and `sense-diff-facing`, there are routines that effect robot movement, change of heading, and acquisition of sensor information.

In a second experiment, we tested whether the *Climber* class of movement strategies were more commonly prevalent with increasing lengths of runs. 25 runs were conducted with the same conditions as in the initial set of runs, except that the runs lasted for 5,000,000 updates instead of 500,000.

III. RESULTS

A. Experiment 1

To provide a basis for comparison, a hand-coded Avidian that used a random walk was injected in a run with mutation turned off. There was no apparent correlation of movement of the random walker with increasing resources (see Figure 2).




Fig. 2. Visits to grid cells by random walker, with 2D projection of resource distribution. The scale indicates the number of visits made to each grid cell. The resource distribution is shown in 2D and is not to scale.

Eight generally distinctive movement strategies emerged in our 100 runs (see Table I). The three movement strategies that were most often prevalent at the ends of runs were *Cockroach*, *Drunkard*, and *Climber* (see Figure 3). The five other movement strategies are represented in Figure 4. The classes of movement strategies were distinguished by

patterns of visits to grid cells by populations of Avidians. The *Cockroach* movement strategy is so named because the Avidians using this strategy move primarily along the grid boundaries and via diagonals across the grid interior. Populations in which *Drunkard* is prevalent show visits to grid cells that, in aggregate, are more common in the vicinity of the resource peak than elsewhere in the grid. The *Climber* movement strategy shows a sharp peak in grid cell visits at the resource peak. The *Column* movement strategy is recognizable from the aggregate visits to grid cells because most visits occur in grid cells of a single column. The *Row* movement strategy similarly is recognizable because most visits occur in grid cells of a single row. The *Cake* movement strategy shows a plateau of visits centered on the resource peak. The *Square* movement strategy shows a similar plateau of visits, but instead of a circular boundary to the plateau, it shows distinct corners. The *Shark* movement strategy shows a characteristic triangle-shaped peak of visits along one side of the grid.

Every evolutionary run produced at least one recognizable movement strategy. Because each run began with an Avidian whose only capability was self-replication, initially what movement occurred showed no such recognizable movement strategy. Almost all runs resulted in the successive evolution of two or more movement strategies (e.g., Figure 7). The first plot in the figure shows the start state, where the `move` instruction is sometimes incorporated into Avidian genomes, but no recognizable pattern of movement has yet resulted. The higher numbers of visits to grid cells on the left-most boundary is explained by a regularity in Avida, that the facing associated with grid cells is initialized in these runs to point to the left. Avidians encountering a grid boundary will have their facing turned along the edge of the boundary half of the time when assigned a random legal facing. The second plot shows the emergence of the first recognizable movement strategy in this run, *Cockroach*. This occurs prior to update 15,000 in this particular run. The second evolved movement strategy, *Drunkard*, appears in the fourth plot, and became prevalent in the population prior to update 25,000. The third evolved movement strategy, *Column*, appears in the tenth plot (or second plot in the third row), and became prevalent prior to update 265,000 in the run. The fourth and final evolved movement strategy, *Climber*, appears in the fourteenth plot (or second plot in the fourth row), having become prevalent prior to update 375,000 in the run. The preceding plot clearly shows a mix of *Drunkard* and *Climber* Avidians contributing to the grid cell visits.

Different runs produced different implementations of the strategies. For example, *Climbers* evolved in 12 runs out of 100, but some used `tumble` very frequently and others used it sparingly, thus showing twisty paths or long straight-line paths, respectively, in regions without large differences in resource availability (see Figure 5). Thus the collective evolved behaviors display the effects of both convergence due to the shared environment and the contingency of the mutations that produced the variations upon which selection

TABLE I
MOVEMENT STRATEGIES SEEN IN EXPERIMENT 1, 500,000 UPDATES

Name	Description	Appearance in Runs	Prevalent at End of Runs	Percent
Cockroach	Moves around boundaries and diagonals	93	24	24%
Drunkard	Biased random walker	80	57	57%
Climber	Approximates gradient ascent	12	12	12%
Row	Favors a particular row	23	0	0%
Column	Favors a particular column	19	3	3%
Cake	Circular plateau of visits centered on resource peak	3	3	3%
Square	Rectangular plateau centered on resource peak	1	1	1%
Shark	Peak visits shows as triangle at side	1	0	0%

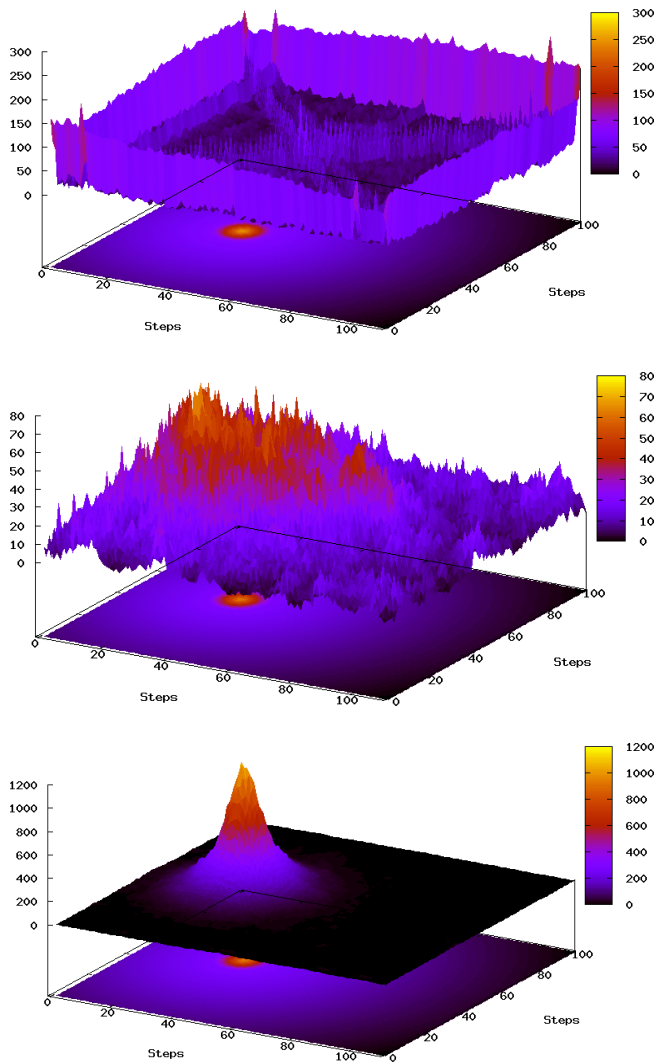


Fig. 3. Most common types of evolved movement strategies. Visits to grid cells over 5,000 updates plotted for *Cockroach*, *Drunkard*, and *Climber* movement strategies, yielding a snapshot of space-time dynamics in the populations.




Fig. 4. Visits to grid cells over 5,000 updates plotted for five other movement strategies: *Cake*, *Square*, *Row*, *Column*, and *Shark*.

acted.

Another characteristic of intelligence is the ability to generalize a strategy and use it in a new situation. An evolved *Climber* could successfully move to resource peaks set in new locations (see Figure 6).

Transitions between states representing the eight observed classes of movement strategies plus a start state were tallied for every 5,000 updates in each run up to update 495,000 (see Table II). Each zero entry in the table indicates that there were no observed instances of a population transition from the movement strategy or start state to the movement strategy indicated at the top of the column. Examining the column headed by a particular movement strategy for non-zero entries reveals which movement strategies were precursors for it. The probability of self-to-self transition is a good indicator of stability. The *Climber* movement strategy is only a precursor for itself, as is the *Square* movement strategy. All other movement strategies have been observed to be precursors for a different movement strategy. Five

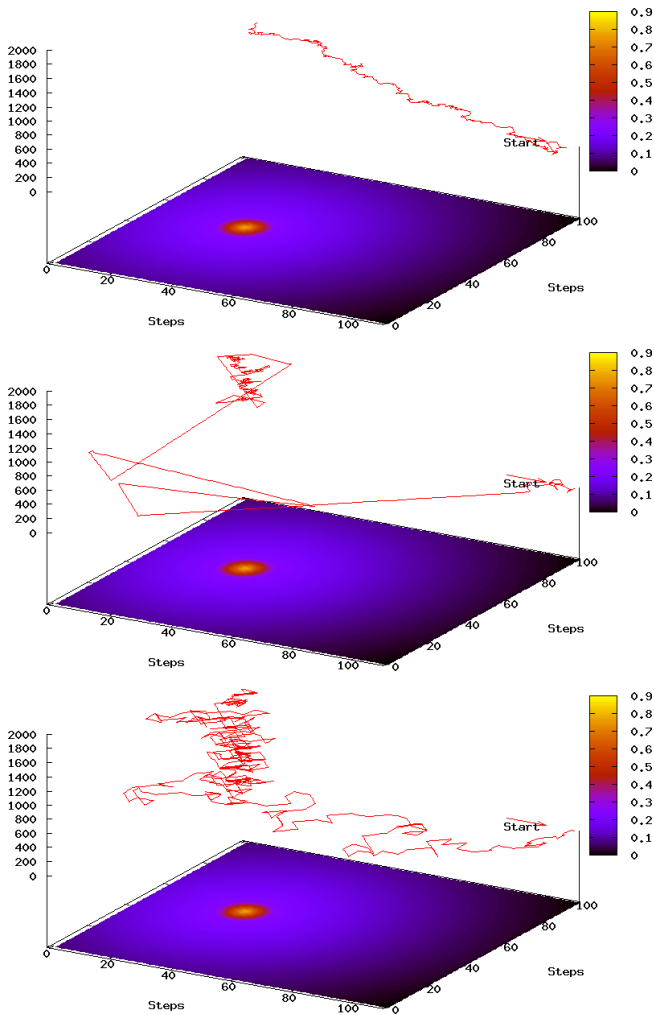


Fig. 5. Diversity of *Climber* Avidians shown with three representative examples run for 2,000 updates. The top Avidian takes almost 2000 updates to reach the resource peak and is thus relatively inefficient, the bottom one is most efficient of these, and the middle Avidian demonstrates an intermediate efficiency.

other movement strategies have been observed as precursors to *Drunkard*. Four other movement strategies have been observed as precursors to *Cockroach*.

B. Robotic embodiment

Preliminary results when translating high performance Avidians to the iRobot Create platform showed the same recognizable differences in behavior between classes of movement strategies. The classes of movement strategies broadly differ in the proportion of `move` instructions to `tumble` instructions, and that difference is appreciable in robot behavior.

C. Experiment 2

Results from our set of 25 runs lasting 5,000,000 updates are summarized in Table III. The *Climber* class of movement strategies evolved to be prevalent in 20 of these runs, or 80% of the final outcomes.




Fig. 6. The same *Climber* Avidian is used in Avida runs with randomly-placed resource peaks. Three representative examples of such runs are shown. The generality of the *Climber* movement strategy is seen in the way the same Avidian finds the randomly-placed resource peaks.

D. Monte Carlo simulations

Using the Monte Carlo methods described earlier, simulated sets of runs produced the following results. At the 5,000,000 update mark, the tally of *Climber* movement strategy as the end state in 1000 simulated runs was 742, or 74.2%. This compares fairly well to our result of 80% of our actual runs in our second experiment showing the *Climber* movement strategy as the final prevalent movement strategy. Dividing the 1,000 simulated runs into sets of 25 reveals that 14 out of 40 sets had 20 or more runs ending with the *Climber* movement strategy.

IV. DISCUSSION

In our first experiment, the three principal evolved movement strategies were *Cockroach*, *Drunkard*, and *Climber* (Figure 3). These span a range of uses of environmental information. *Cockroach* appears to be an

TABLE II
TRANSITION PROBABILITY BETWEEN STATES SEEN IN EXPERIMENT 1, 500,000 UPDATES

Initial state	Start	Cake	Climber	Cockroach	Column	Drunkard	Row	Shark	Square	Total
Start	0.68	0.00	0.00	0.29	0.00	0.036	0.00	0.00	0.00	308
Cake	0.00	0.99	0.00	0.00	0.00	0.0075	0.00	0.00	0.00	133
Climber	0.00	0.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00	425
Cockroach	0.00	0.0003	0.0003	0.96	0.012	0.020	0.0096	0.00	0.00	3835
Column	0.00	0.00	0.0057	0.26	0.34	0.36	0.029	0.00	0.00	174
Drunkard	0.00	0.0006	0.0021	0.0035	0.013	0.98	0.0017	0.0002	0.00	4823
Row	0.00	0.00	0.00	0.27	0.079	0.047	0.61	0.00	0.00	127
Shark	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.98	0.014	71
Square	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	1.00	4
Count	208	136	437	3857	178	4881	127	71	5	9900

TABLE III
OUTCOMES OF LONG RUNS

Name	Prevalent at end	Percent
Climber	20	80%
Drunkard	3	12%
Cake	2	8%

application of blind search that exploits the fact that the environment here is a bounded grid. Movement around the perimeter plus occasional crossings of the interior space on the long diagonals mean that an Avidian employing this strategy will approach within about 20% of the length of the longest side to even the most inaccessibly-placed resource peak. Given the monotonically-decreasing distribution of our resource and coupled with the random initial placement of Avidians on birth, *Cockroach* provides a better movement strategy than a localized random walk. If intelligence is taken to be the increased frequency of emission of adaptive behavior under novel stimuli, as is seen in studies of animal behavior, *Cockroach* either does not qualify as such, since all stimuli yield the same behavior, or may be seen as a small relative improvement on a random walk, since it does exploit the conditions implicit in a bounded grid for movement. *Drunkard* is not simply a random walk; it is instead a biased random walk that in aggregate movement of the population visits the region around the peak resource somewhat more often than other parts of the grid. Information use is present, but not efficient for *Drunkard*, representing an improvement over the *Cockroach* movement strategy. *Climber* moves effectively toward the resource peak using the information in the resource gradient and then remains in the vicinity of the resource peak. The test of generality shows that the *Climber* programs are not dependent upon the configuration used in the evolution runs, but rather are capable of locating a randomly-placed resource peak. (Generality is simplest to uncover for the *Climber* movement strategy, and will require substantial additional computer time to test for other movement strategies.) Goal-seeking behavior, such as self-direction towards a target, is one common diagnostic character of intelligence and simple forms of such behavior

have evolved in these populations.

Transition probabilities tell us useful information about convergence and stability of movement strategies. Notably, our class of gradient ascent programs, *Climber*, is highly stable: in all instances where the population once has *Climber* as the prevalent movement strategy, it remains that way thereafter. The transition table (Table II) reinforces the insight that evolutionary processes *contingently* build upon existing capabilities to derive new ones. Various movement strategies only evolve given the prior evolution of other, precursor, movement strategies. The likelihood of evolving a dependent movement strategy has to do with both the frequency with which the precursor occurs, and with the inherent stability of the precursor movement strategy. The more stable a precursor strategy is, the less frequent is the evolution of the dependent movement strategy. Precursor states may provide modules or features necessary for co-option to evolve new functionality [14]. The *Climber* movement strategy was observed to be *reachable*, directly or indirectly, from all other movement strategies save *Shark* and *Square*, and those were rare within our dataset. The *Climber* movement strategy, though, was only reached directly from populations where the *Cockroach*, *Drunkard*, or *Column* movement strategies were prevalent. In our simple system, the common pattern of progression in movement strategies is the early evolution of *Cockroach*, followed by evolution of *Drunkard*, and in some cases the evolution of *Climber* follows. Our second experiment using much longer runs shows that the *Climber* movement strategy becomes the most commonly prevalent movement strategy at the end of runs, given a longer time for evolution to occur.

There are some interesting implications of the results of our Monte Carlo simulations based upon the transition probabilities derived from our Experiment 1 results. These were sufficient to generate an expectation compatible with the results of our Experiment 2 with respect to the probability that the *Climber* movement strategy would be prevalent at the end of runs. This indicates a certain degree of predictability of outcomes when the analysis is conducted at the level of classes of behavior. It also implies that the underlying contingent dynamics seen in shorter runs are adequate to explain the results of our longer runs, and that longer runs




Fig. 7. Each graph above represents total Avidian visits to grid cells for 5,000 updates in a run for points near transitions between movement strategies. The progression shows the initial start state, then *Cockroach*, then *Drunkard*, one period of *Column*, and ends with *Climber* as the prevalent movement strategy.

have not introduced a qualitative difference in outcomes compared to the shorter runs.

It is characteristic of evolutionary processes to accrue and expend variation in a population. Within our experiments, we saw variation at two levels. Examination of individual Avidian programs from within the most common classes of movement strategies revealed one level of variation, as Avidians within particular movement classes differed from one another. Within the *Climber* class of movement strategies, there were notable differences in the efficiency of individual Avidian programs (Figure 5). The higher level of variation was seen in the diversity of the classes of movement strategies. We observed a total of eight movement strategies with recognizably different patterns of movement based on aggregate population behavior.

The evolved programs of Avidians are examples of *effective methods*, a finite series of steps that when carried out deliver a particular outcome. What is rewarded within Avida is always simply the accrual of higher merit. We can recognize classes of behaviors evolved by Avidians and cast those in other terms, as in the case of the class of *Climber* movement strategies that approximate gradient ascent approaches.

Our results exemplify the evolution of *emergent behavior*, with a multiplicity of solutions produced in response to a simple environmental condition, elaborated by the context of a population of competitors. The phenomenon of *satisficing* is seen in our results, with a non-optimal movement strategy, *Drunkard*, being the most common strategy seen at the end of runs in our first experiment, and still seen even in the far longer runs of our second experiment. Satisficing describes cases where non-optimal adaptations persist in evolving populations, even when we have an expectation that better adaptations may have originated as rare traits in those populations. We propose a scenario that explains the appearance of satisficing between movement strategies

in our experiments. An important consideration here is that success in evolutionary terms depends not only upon some figure of merit for a specific task of interest, but upon the relationship of an individual to the general performance of other individuals in the population. With the pattern of the successive emergence of different movement strategies seen in our experiments, it is apparent that for a novel movement strategy to displace a previously prevalent strategy, it must provide an appreciable advantage over that strategy, and it must not be lost due to drift. This need not necessarily be an advantage based on the task of interest, but could also be premised upon such factors as reducing gestation time. Those “non-task” factors will also be acted upon within the sub-population that performs the current prevalent movement strategy. From the perspective of an Avidian, there is no distinction between task and non-task. In our evolution runs, the population size was limited to 200 organisms or less. While not tiny, this is nonetheless a fairly small population size where appreciable effects of drift might be seen. As in biological evolution, those solutions that lead to greater representation of similar organisms in future generations displace others. This puts a novel competitive movement strategy at a disadvantage in that the number of individuals initially performing it is small, and thus may be lost due to drift, and that improvements in non-task factors are more likely to occur in the larger sub-population that is using the current prevalent movement strategy. The conditions for switching a population between movement strategies may at times require a combination of improvements in task performance and non-task factors, plus an additional portion of luck to evade winnowing by drift when individuals performing the new movement strategy are rare. These considerations, we believe, provide an explanation for the phenomenon of satisficing we observed that is based in known population genetics and dynamics.

A. Future Work

Even within the context of a simple environment composed of one resource peak, there was considerable complexity seen in the evolved outcomes. There are a number of analyses requiring extensive additional runs that could be conducted, but are beyond the scope of this paper.

There are several further questions that we are contemplating. Would expansion of the resource space to three dimensions alter the basic movement strategies seen or their common sequence of occurrence? The different outcomes of likelihood of visiting nodes on a lattice at two and three dimensions [17] suggests that movement strategies or relationships between transitions may differ in the two cases. This involves stepping away from the current two-dimensional cell layout used in Avida, and replacing it with a flexible topology-management system that will provide both a continuous and three-dimensional space for Avidians to move in. Another question is whether evolution can find effective methods in the case where a spatially-defined resource also shifts position over time during a run, which should

also permit investigations into evolved behavior concerning periodicity in resources and foraging under uncertainty.

Another question is how a more complex environment would change the evolution of movement strategies, specifically in the case where both an *appetitive* resource distribution and an *aversive* resource distribution exist, as in having both “food” to be located and “poison” to be avoided within the environment.

We also want to explore competitions between our evolved Avidian programs. This should help clarify what features of evolved programs help make them successful, whether those have to do explicitly with resource exploitation or not.

V. CONCLUSIONS

In testing the capability of evolutionary computation to produce effective methods utilizing movement strategies to intelligently exploit spatially-distributed resources, our results show that such strategies do emerge and that in about 12% of shorter runs and in 80% of longer runs the final movement strategy used by the majority of the population at the end of the run is in the class of optimal response for our environment, that of gradient ascent.

The pattern of emergence of movement strategies was not strictly fixed, but in most runs there were multiple movement strategies employed by most of the population over the time course of the run. Contingency played a role, as evidenced by the multiplicity of movement strategies, diversity of individual implementations, and the variety of movement strategies employed by most of the population at the end of our runs.

Given that our environment utilized an extremely simple spatial resource distribution, that of a resource that monotonically decreases with increasing distance from the resource peak, the observed set of eight movement strategy classes and apparent diversity of evolved programs within those classes indicates that the degree of complexity of evolved responses is not tightly constrained by the complexity – or simplicity – of the environment. The degree of interaction between Avidians during the evolution runs was small, so any contribution of inter-Avidian interactions to the evolved complexity should likewise be small.

Satisficing broadly describes phenomena where non-optimal solutions that are “good enough” are retained. Our results are consistent with an interpretation of satisficing in that the optimal class of response, a *Climber* movement strategy approximation to gradient ascent, was only seen in a relatively small fraction of our shorter runs, rising to a larger fraction of longer runs when other strategies are no longer satisfied given an environment with stronger competitors. The most common movement strategy used by the majority of a population at the end of shorter runs was instead *Drunkard*, a sub-optimal biased random walk. The *Drunkard* and *Cake* movement strategies were the only non-

Climber movement strategies seen at the end of longer runs, indicating that for some conditions those movement strategies remain “good enough.”

VI. ACKNOWLEDGMENTS

We wish to thank our colleagues of MSU’s Digital Evolution Lab: Phil McKinley, Fred Dyer, Richard Lenski, Ben Beckmann, Jeff Barrick, Jacob Walker, and Brian Baer for their advice and assistance. This research is supported by a grant Emerging Intelligence: Contingency, Convergence and Constraints in the Evolution of Intelligent Behavior from the Cambridge Templeton Consortium and National Science Foundation Grant CCF-0643952.

REFERENCES

- [1] J. Hopfield and D. Tank, ““neural” computation of decisions in optimization problems,” *Biological Cybernetics*, vol. 52, no. 3, pp. 141–152, 1985.
- [2] J. Roese, K. Risenhoover, and L. Folsie, “Habitat heterogeneity and foraging efficiency: An individual-based model,” *Ecological Modelling*, vol. 57, no. 1, pp. 133–143, 1991.
- [3] S. Forrest and T. Jones, “Modeling Complex Adaptive Systems with Echo,” *Complex Systems: Mechanism of Adaptation*, 1994.
- [4] G. Booth, “Gecko: A Continuous 2D World for Ecological Modeling,” *Artificial Life*, vol. 3, no. 3, pp. 147–163, 1997.
- [5] S. Yamada, “Evolutionary behavior learning for action-based environment modeling by a mobile robot,” *Applied Soft Computing*, vol. 5, no. 2, pp. 245–257, JAN 2005.
- [6] K. Matsushita, H. Yokoi, and T. Arai, “Pseudo-passive dynamic walkers designed by coupled evolution of the controller and morphology,” *Robotics and Autonomous Systems*, vol. 54, no. 8, pp. 674–685, 2006.
- [7] S. Mahdavi and P. Bentley, “Innately adaptive robotics through embodied evolution,” *Autonomous Robots*, vol. 20, no. 2, pp. 149–163, MAR 2006.
- [8] A. Montebelli, C. Herrera, and T. Ziemke, “On cognition as dynamical coupling: an analysis of behavioral attractor dynamics,” *Adaptive Behavior*, vol. 16, no. 2-3, pp. 182–195, 2008.
- [9] P. Nordin and W. Banzhaf, “Genetic programming controlling a miniature robot,” in *Working Notes for the AAAI Symposium on Genetic Programming*, 1995, pp. 61–67.
- [10] C. Ofria and C. O. Wilke, “Avida: a software platform for research in computational evolutionary biology,” in *Artificial Life 10*, 1994, pp. 191–229.
- [11] C. Adami, C. A. Ofria, and T. C. Collier, “Evolution of biological complexity,” *Proceedings of the National Academy of Science*, vol. 97, pp. 4463–4468, 2000.
- [12] C. Ofria, C. Adami, and T. C. Collier, “Design of evolvable computer languages,” *IEEE Transactions in Evolutionary Computation*, vol. 17, pp. 528–532, 2002.
- [13] D. C. Dennett, “The new replicators,” in *Encyclopedia of Evolution*, M. Pagel, Ed. New York: Oxford University Press, 2002, pp. E83–E92.
- [14] R. E. Lenski, C. Ofria, R. T. Pennock, and C. Adami, “The evolutionary origin of complex features,” *Nature*, vol. 423, pp. 139–144, 2003.
- [15] L. M. Grabowski, W. R. Elsberry, C. Ofria, and R. T. Pennock, “On the evolution of motility and intelligent tactic response,” in *GECCO ’08: Proceedings of the 10th annual conference on Genetic and evolutionary computation*, July 2008.
- [16] B. Beckmann, L. M. Grabowski, P. McKinley, and C. Ofria, “An autonomic software development methodology based on Darwinian evolution,” in *The 5th IEEE International Conference on Autonomic Computing*, Chicago, June 2–6 2008.
- [17] W. H. McCrea and F. J. W. Whipple, “Random paths in two and three dimensions,” *Proc. Roy. Soc. Edinburgh*, vol. 60, pp. 281–298, 1940.