

A robot trace-maker: modeling the fossil evidence of early invertebrate behavior

Tony J. Prescott and Carl Ibbotson
Department of Psychology
University of Sheffield
Western Bank
Sheffield S10 2TP

email: t.j.prescott@sheffield.ac.uk

Abstract

The study of trace fossils, the fossilized remains of animal behavior, reveals interesting parallels with recent research in behavior-based robotics. This article reports robot simulations of the meandering foraging trails left by early invertebrates which demonstrate that such trails can be generated by mechanisms similar to those used for robot wall-following. We conclude with the suggestion that the capacity for intelligent behavior shown by many behavior-based robots is similar to that of animals of the late Precambrian and early Cambrian periods approximately 530-565 million years ago.

keywords: trace fossils; behavior-based robots; foraging behavior.

1 Introduction

Biology is widely regarded as an important source of inspiration for robotics (e.g. [20]). Animals are seen as offering working examples of robust, embedded autonomous agents, and their neural circuitry, and sensor and motor structures, are viewed as providing models for designing similar components for robots. Ideas are also beginning to flow from robotics to biology. For instance, the control architecture for a six-legged robot developed by Brooks [6] has provided a useful metaphor for understanding the functional architecture of insect nervous systems [1]; other robots, of various designs, are being used to embody and evaluate theoretical models developed in biology (e.g. [2, 17, 45]). All robots, like all animals, must physically interact with their surroundings. The study of robot models of animal behavior can therefore make an important contribution to understanding how the nervous system, body, and environment interact to generate adaptive behavior [9].

A further premise of this article is that robotics could benefit from taking a closer interest in evolutionary history—the study of the phylogenetic relationships between animals and the nature of evolutionary change from the first animal forms to those of the present day. Valuable insights for robotics should be gained if we can understand how complex biological control systems were derived from simpler ones—a question that can only be answered by investigating the evolution of neural circuits and of animal behavior. Of course, brains

and behavior don't make good fossils. This gives investigators of neural and behavioral evolution the difficult task of drawing inferences about ancestral forms from clues found in comparative and developmental studies (e.g. [8, 16, 30, 33]). Only occasionally does the fossil record cast a direct light on the behavior of extinct organisms. In this article, however, we are concerned with fossil evidence of exactly this sort which has provided important insights into the behavioral competencies of some of the first bilateral invertebrates. Perhaps surprisingly, we will see that the nervous systems of these creatures generated behavior with remarkable similarities to that of some recent, simple, 'behavior-based' robots. The study of such robots, we will suggest, could also help to illuminate the biological understanding of these ancient animals.

2 Trace Fossils

The common ancestor of all modern, bilaterally symmetric, metazoans (multicellular animals) was probably a roundish worm that lived in a shallow marine environment during the Vendian period (544 to 565 million years ago) of the late Precambrian [43]. Such creatures, and many of the organisms (worms, mollusks, and arthropods) they gave rise to, left no actual fossils as they had virtually no hard body-parts. Fortunately, however, these animals did leave a fossil record of sorts—the tracks, trails, and burrows that some early invertebrates made in the sediment have been preserved forming what are now called trace fossils. Some of the commonest forms of trace fossil record foraging trails left on, or just below, the sea bed. The Earliest traces reflect simple

‘scribbling’ behaviors, with tracks that often cross themselves, and indicate relatively crude foraging strategies. By the end of the Vendian, however, more regular trails appear that form spirals or ‘meanders’ that loop back on themselves without crossing. The early Cambrian (530-544 million years ago) sees a further increase in the abundance, complexity and diversity of trace fossils with trails and burrows that penetrate deeper into the substrate and possess multiple levels and branches [13]. The changes that occurred over these time-spans reflect three important developments: an increase in the diversity of animals, improvements in burrowing capabilities, and most importantly, an increase in the complexity of neural circuits. The beginning of the Cambrian period is, of course, also recognized as marking the origin of the contemporary metazoan phyla. The Cambrian ‘explosion’ saw the rapid emergence, over the course of ten to twenty million years, of a diversity of body forms equipped with relatively complex sensory and nervous systems. Trace fossils therefore represent a primary source of insights into the sequence of evolutionary events that anticipated the appearance of the modern fauna [33, 43].

3 Computer models of trace fossils

In attempting to infer the behavioral capabilities of the ancient animals that left fossil traces, it seems reasonable to seek the simplest mechanisms that will reproduce the observed patterns. Following Braitenberg’s [3] advice that “when we analyze a mechanism we tend to overestimate its complexity” the methodology of synthetic psychology—building model systems that generate similar behaviors—seems an appropriate strategy. In fact, synthetic approaches were applied some time ago to the understanding of trace fossil behavior; the computer simulations of Raup and Seilacher [34], first published in 1969, standing out as an early, and rarely acknowledged, example of what might now be termed Artificial Life.

The most consistent fossilized foraging patterns were formed in areas of the sea bed with an even distribution of food particles in the sediment. This environment favors compact trails with maximal coverage and minimal recrossing of existing tracks. Compared with straight-line movement a meandering pattern also helps to keep the animal within its preferred environment and may reduce the likelihood of encountering competitors. Raup and Seilacher, following Richter [35], based their models of these trace fossils on a combination of simple reactive behaviors: *thigmotaxis* that makes the animal stay close to previously formed tracks; *phobotaxis* that causes it to avoid crossing existing tracks; and *strophotaxis* that causes it to make 180° turns at various intervals. A fourth behavior causes the

animal to move forward when the conditions of the other three are not met. The computer simulations performed by Raup and Seilacher demonstrated that the interaction between these behaviors is sufficient to generate the coiled meandering patterns characteristic of many foraging trails. Figure 1, taken from [34], show some typical meanders generated by their program together with the trace fossils they were designed to emulate. Part 2 of figure 1 has the interesting feature that the thigmotaxis response is particularly weak—following a U-turn the animal takes some time to restore contact with its earlier track. This relatively inefficient foraging behavior, a characteristic of some early fossils, is taken by Raup and Seilacher as evidence that thigmotaxis and phobotaxis are “genetically distinct behavioral reactions”. The foraging meander is therefore seen as an emergent pattern arising from the environment-mediated interaction of a number of distinct behavioral competencies.

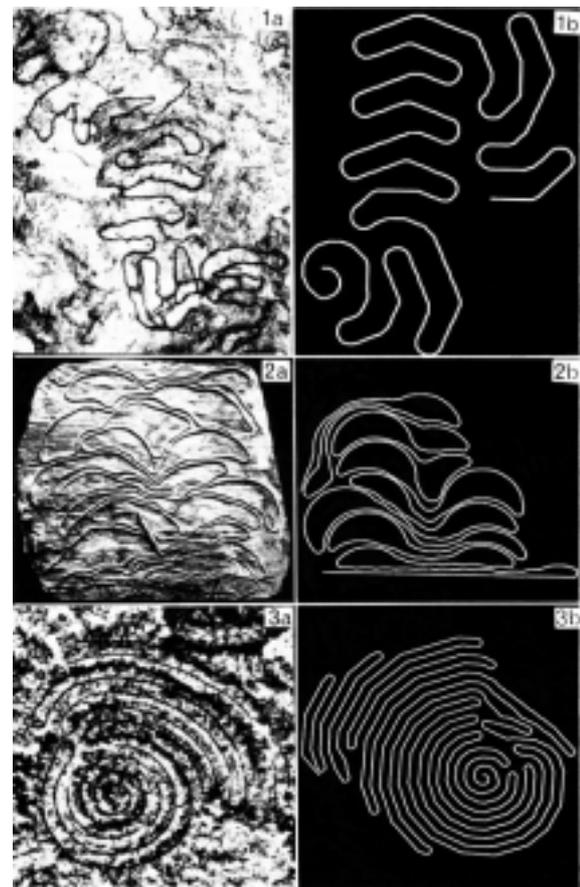


Figure 1. Trace fossil meanders and comparable computer output. From [34]. 1. Meandering trail of a trace fossil of genus *Dictyodora* (Ordovician period). 2. Loose meandering pattern of genus *Helminthoraphe* (Upper Cretaceous period). 3. Burrow of the living beach worm *Paranois fulgens*.

Seilacher and coworkers [18, 34, 37-39] suggested a number of further models for different types of foraging trace, and proposed that some of the variation between fossil traces, and some of the changes that occur over the course of evolution, could be modeled by manipulating key parameters of the various component reactions. For instance, the turning radius of the animal, the mean distance between tracks, and the relative strengths of phobotaxis and thigmotaxis, can each be varied to generate trails with different characteristic patterns and varying degrees of foraging efficiency.

4 Robot models of trace fossils

The principle of a complex behavior pattern emerging from the competitive interaction of a number of simple reactions is a characteristic that Raup and Seilacher's work clearly shares with the 'behavior-based' robotics approach of Brooks and others (e.g. [5, 7, 11, 21, 27, 29]). Indeed, the simulated meander generator invites comparison to the 'wall-following' mechanisms that have been investigated in a number behavior-based mobile robots (e.g. [11, 28, 31, 41]). For instance, Mataric [28] describes a robust wall-following behavior that emerges from the interaction of three modules. The first, similar thigmotaxis, causes the robot to steer inwards (toward the wall) when sensors indicate that the distance to the wall is above some threshold; the second, similar to phobotaxis, causes the robot to steer outwards when the sensed distance is below a threshold; while, the third, causes the robot to move forward whenever the conditions of the other two behaviors are not met. It is apparent from this description that (apart from the additional strophotaxis component) the meander generator described by Raup and Seilacher is remarkably similar to a robot wall-follower where the object being followed, instead of being a fixed contour, is the trail of disturbed sediment generated by the animal's own movements.

Our combined interest in evolutionary biology and robotics has led us to investigate trace-making behavior in a custom-built robot that generates and follows trails across the laboratory floor. Our aim is to enhance the realism of some aspects of trace fossil modeling (for instance, by introducing the constraints of genuine sensorimotor coordination), and, at the same, to demonstrate the parallels between behavior-based robots and the behavior of early metazoan animals. The sediment feeders we have attempted to model probably used chemical and mechanical sensory systems to detect and follow their tracks and burrows [37]. However, as a first approximation to these mechanisms we have used light sensors to detect a trail of paper which is dispensed by the robot as it moves. The robot trace-maker, shown in figure 2, possesses two motor-

driven wheels, a front non-driven castor (to provide stability), and a motor-driven paper dispenser which releases a constant stream of paper while the robot is moving. Steering is effected by changing the speed of the two driven wheels each of which has five speeds (fast forward, slow forward, stop, slow reverse, and fast reverse). Two light dependent resistors (LDRs) are mounted on each of two lateral arms. The ability to detect and follow the paper trail relies on the different reflectance (under ambient light) of the white paper dispensed by the robot and the dark-colored laboratory floor on which the robot moves. Two infra-red sensors mounted at the front of the robot are used to implement a primitive obstacle detection mechanism. Further details of the robot implementation are given in the Appendix.

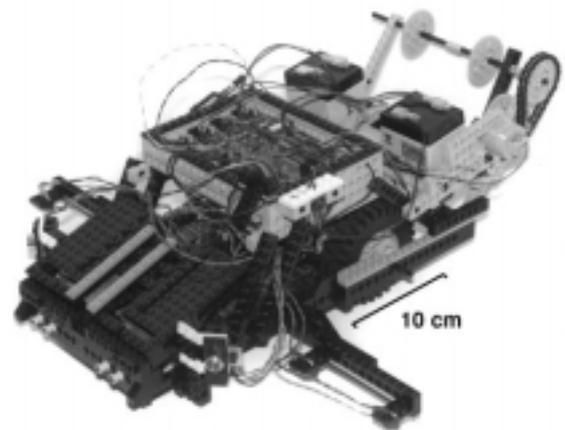


Figure 2. The robot trace-maker. The motorized dispenser on the back of the robot releases a stream of paper when the robot is moving. Two light detectors (LDRs) on each of the side arms measure reflected light from the floor and control the thigmotaxis (toward track) and phobotaxis (away from track) behaviors. Infra-red sensors on the front of the robot support avoidance behavior.

Spiral trails

A spiral trace pattern can be generated, using just one of the two sensor arms, and a combination of *thigmotaxis* (approach existing track), *phobotaxis* (avoid track re-crossing), and a default *advance* behavior. Phobotaxis, interpreted as a turn away from the sensor arm, is triggered when the value of the inner LDR sensor is greater than a threshold set slightly above the reflectance of the dark floor surface. Thigmotaxis, a turn towards the sensor arm, is triggered when the value of the outer LDR sensor is less than a threshold set slightly below the reflectance of the paper dispensed by the robot. In both cases the required turn is achieved by rotating the drive wheel on the outside of the turn at speed 'slow forward' and that on the inside at speed 'slow reverse'. Phobotaxis has priority over thigmotaxis so

that in the rare event that trigger conditions for both behaviors are satisfied simultaneously the former takes precedence (this can happen, for instance, if the paper trail becomes twisted and is therefore uncharacteristically narrow at a given point). When neither of the two taxes are active the *advance* behavior moves the robot forward on a trajectory that drifts inwards towards the sensor arm. This is achieved by driving the inner wheel at speed 'slow forward', and the outer wheel at speed 'fast forward'.

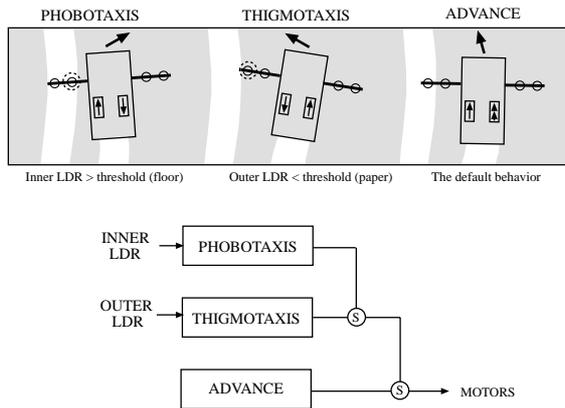


Figure 3. The robot architecture for generating spiral trails combines phobotaxis, thigmotaxis, and an advance behavior. The upper diagrams show the trigger conditions for each behavior and the resulting motor output to the two driven wheels. The lower diagram shows the control system viewed as a 'subsumption' architecture—higher priority behaviors subsuming (s) those below.

Figure 3 illustrates these behaviors and shows the robot control system viewed as a 'subsumption architecture' [5, 7]. In such architectures a high priority behavior can suppress the output of a low priority one and substitute its own (subsumption). Thus, in our architecture, phobotaxis subsumes thigmotaxis which subsumes the advance behavior.



Figure 4. A spiral 'foraging' trail generated by the robot trace-maker.

Figure 4 shows the trail generated when behavior is driven by the left sensor arm. The robot spirals

outwards in an anti-clockwise fashion forming a pattern, known as the *Archimedes spiral*, which is often found in trace fossils. Note, that the start of the spiral needs no particular programming but is simply the outcome of a situation in which there is previous trail to follow. Figure 5 shows a spiraling trace fossil for comparison.



Figure 5. A spiral foraging trail of genus *Spirodesmos* (Mississippian period). Reprinted with permission from [24].

More complex meanders

Strophotaxis, the U-turn behavior, can be effected in a simple manner by transferring control of thigmotaxis from one sensor arm to the other. The angle of turn need not be specified as thigmotaxis on the newly activated side of the robot will cause it to rotate until contact with the trail is regained on that side. A general architecture for generating complex meanders containing such U-turns is illustrated in figure 6.

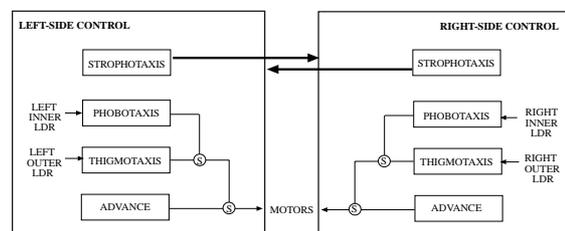


Figure 6. An architecture for generating meandering trails. Only one side of the control system can be active at a given time. The active strophotaxis behavior switches control to the opposite side after a specified time interval.

The *phobotaxis*, *thigmotaxis*, and *advance* modules in the left-side of this control circuit are the same as in the architecture used to generate the spiral; those on the right-side are copies with appropriate connections to the motors and to the right sensor arm. Only one side of the architecture is active at a given time. The strophotaxis module in the active side initiates a U-turn by switching off the control modules on its own side and activating those on the opposite side.

An interesting question concerns how and when the strophotaxis behavior is triggered. Raup and Seilacher [34] have suggested that a worm-like animal might use the length of its body to determine when to U-turn. The length of the lobes of a trace fossil (the sections between U-turns) might be equal to the length of the animal's body, implying that the trigger for a turn occurred when the tail of the animal emerged from the previous bend. Of course, many foraging patterns contain lobes of varying length. To account for these variations Raup and Seilacher suggested that shorter lobes might be generated when an animal encountered an obstruction that disrupted its forward motion and precipitated an early turn. Longer than average lobes may have arisen when the tail of the animal detected a minor bend in the lobe, which is falsely identified as a U-turn, causing it to continue in a forward direction. It is difficult to evaluate this "tail-straightening" hypothesis given that the identity, and many of the morphological characteristics, of the early trace fossil generators are unknown; however, it does seem to provide a plausible theory for strophotaxis behavior.

The robot trace-maker is obviously neither long-bodied nor articulated hence the length of the lobes cannot be easily determined by reference to the robot's body. Instead we have chosen to implement strophotaxis using a *time-out*—a mechanism frequently employed in behavior-based robotics—where an event is triggered after a fixed time interval has elapsed. Hence, at the start of the meander, or when control switches from one side of the robot to the other, a counter is initialized in the newly active strophotaxis module. When the prescribed interval has elapsed this module initiates a strophotaxis event switching control to the opposite side.

Some robot traces generated using time-outs for strophotaxis are shown in figures 7 and 8. In making these trails we have employed two variants on the meandering architecture described above.

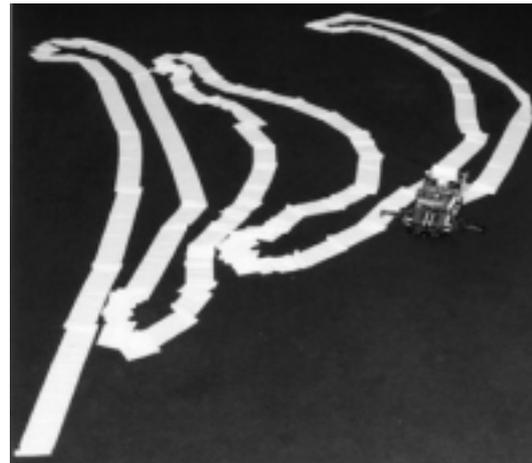
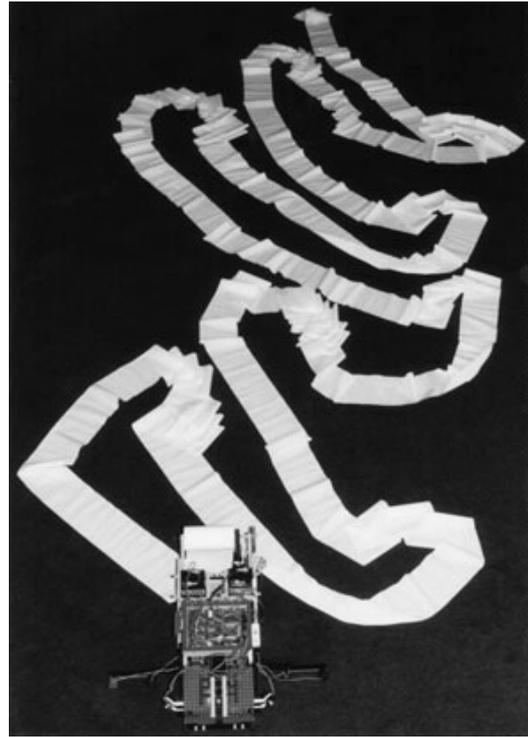


Figure 7. 'Weak' robot meanders. Thigmotaxis is employed to locate the trail immediately following a strophotaxis event but is otherwise disengaged. Strophotaxis events occur every eight seconds in the upper image and every thirty seconds in lower image.



Figure 8. A 'strong' robot meander formed by continuous monitoring of the distance to the previous lobe. Strophotaxis occurs every thirty seconds.

In figure 7 the robot's trajectory is determined by a relatively 'weak' meandering architecture. Here, thigmotaxis is employed to locate the trail immediately following each strophotaxis event but is otherwise disengaged. Under the control of this architecture the robot makes U-turns and avoids crossing its previous path, but does not actively attempt to track the paper trail. The consequence of employing this form of control can be seen in both of the meanders shown which differ only in the length of the interval between strophotaxis events (eight seconds in upper image, thirty seconds in the lower image).

The robot generates a moderately compact meander, however, it sometimes drifts off course a little after regaining contact with the trail following strophotaxis. The result is a slightly inefficient 'foraging' pattern in which some areas are left unvisited. This strategy may be similar to that used by the worm-like animal that generated the trace fossil shown in figure 9. Note the distinctive kinks in the foraging trail where the animal 'bounces' off its previous trail following each turn. It seems unlikely that this animal was continuously monitoring the distance to the previous trail between turns.

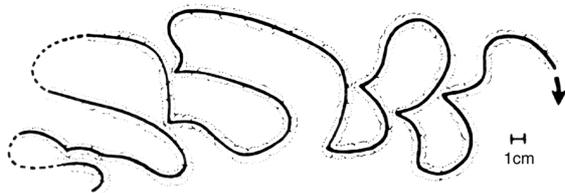


Figure 9. A trace fossil of genus *Nereites* (Cambrian period). Reprinted from [18] with permission of the Scandinavian University Press. The animal that made this trace does not appear to monitor the distance to the previous lobe between strophotaxis events.

Figure 8 shows a more tightly coiled robot trail generated by a 'strong' meandering architecture. In this example, the robot's thigmotaxis and phototaxis behaviors were in continuous operation between U-turns (as they were in the architecture used to generate the spiral in figure 4). The 'push-pull' control of these behaviors ensures that the trajectory of the robot clings closely to the contour of its previous trail generating a more compact pattern. In the example shown, strophotaxis occurs every thirty seconds. However, by adjusting this value, radically different meandering patterns can be generated (settings of around eight seconds or sixty seconds would generate meanders more similar to parts 1 and 3 of Figure 1, respectively).

It is interesting to note that the meander in figure 8 starts out as a spiral pattern. Similar starter spirals have been observed in a variety of different fossil meanders generated by sediment feeders [39], a clear example of which is shown in figure 10 (starter spirals can also be seen in parts 1 and 3 of figure 1). It is evident from the computer and robot simulations of these traces that the same underlying mechanisms can account for both the starter spiral and the meandering components of a trail.



Figure 10. A trace fossil with a starter spiral, genus *Taphrelminthopsis* (Lower Tertiary period). Reprinted from [18] with permission of the Scandinavian University Press.

Finally, in figure 11, we illustrate the addition of a further layer of control to the robot architecture. In this example the infra-red sensors at the front of the robot were used to detect the distance to any nearby obstacles and, when necessary, trigger a high-priority *avoid* module which overrides the meandering behavior. In the example shown, the avoid module was combined with the weaker of the two meander architectures. The intervention of an obstacle avoidance behavior was cited by Seilacher as the possible cause of some of the shorter lobes observed in fossil meanders. It therefore seems likely that the activity of these animals was also controlled by a hierarchy of behavioral competencies.



Figure 11. A robot trail generated by a combination of meandering and obstacle-avoidance behaviors.

Comparisons with computer simulations

There are a number of differences between our own robot simulations and the computer models

investigated by Raup and Seilacher that deserve further comment.

First, to create some variation in their simulated meanders, Raup and Seilacher used a random number generator to determine the lobe length between strophotaxis events (this was based around a mean value conceived to be the body length of the trace-maker). The robot simulation, however, works with deterministic rules and fixed timeouts. Variation, within and between trails in our robot experiments arise because minor disturbances in the environment, or in the motor system, tend to be exaggerated by the behavior-generating mechanisms. So, for example, whilst following some lobes, the robot may have had to engage its phobotaxis behavior more often than on others, therefore not travelling quite as far in the time periods between strophotaxis. Alternatively, the robot may get temporarily stuck in a form of ‘dithering’ behavior—alternating between thigmo- and phobotaxis—while following a particularly uneven previous trail. Variations in the roughness of the carpeted floor, or in battery power levels, can cause the robot wheels to go either faster or slower, while small differences in surface reflectance or in incident light-levels can effect the triggering conditions for the various behaviors causing the robot to over- or under-shoot a turn. Unplanned interactions such as these shape the robot’s activity generating a natural variation in meandering behavior not dissimilar to that observed in genuine trace fossils.

A second difference is that in the computer simulations of Raup and Seilacher each strophotaxis event consists of an explicit turn through a fixed angle (usually 180°). In the robot, however, turns are made under the control of thigmotaxis, and are therefore incremental and terminated by feedback. Both procedures generate fossil-like meandering patterns, however, the self-correcting nature of the robot’s turning mechanism suggests that it may produce more efficient meanders given the wide-ranging sources of natural variation just described. A second advantage is that no additional mechanism (to that already available for thigmotaxis) is presupposed for monitoring the angle of turn. A more detailed comparison between robot and fossil meandering patterns than has been performed here could serve to distinguish between these alternative hypotheses of interoceptively and exteroceptively regulated U-turn behavior (a suitable methodology for such an analysis is described in [23]).

The present work could also be extended to model some more complex meandering patterns. For example, figure 12 shows a number of trace fossils from the genus *Cosmoraphe*. This genus which is

characterized by the presence of small, second-order undulations within the larger meanders, presents an interesting challenge for future modeling. Many other complex trace fossils, however, are three dimensional structures and therefore beyond the scope of our current robot trace-maker.



Figure 12. Various trace fossil species of genus *Cosmoraphe*. Trails belonging to this genus are characterized by the presence of small, second-order undulations within larger first-order meanders. Adapted from [38] with permission of the Liverpool Geological Society.

A point emphasized by Seilacher [39] is that the evolution of foraging behavior can occur through changes in the morphology of the animal as much as through improvements to neural control mechanisms. Our robot is easily reconfigurable, which allows experiments using different sensor positions, turning radius, and so on. In practice, however, we have found that there is only limited scope for variability without radical redesign of the robot’s structure. Changes to one aspect of the morphology often have unforeseen and undesirable effects on the robot’s behavior which can only be put right by rebuilding other aspects of the robot configuration. Indeed, there appears to be a web of constraints here such that the different morphological ‘parameters’ are not really independent.

5 Related Work

Seilacher has described four further categories of trace fossils in addition to the class of foraging trails we have been largely concerned with here. These other categories cover crawling tracks, feeding burrows, resting tracks, and permanent dwellings or burrows. Many fossils provide clues only with respect to a small fraction of the organism’s behavioral repertoire. Feeding and dwelling burrows, however, can (like foraging trails) give a more complete insight into the lifestyle of the trace-maker. For instance, the meander pattern shown in part 3 of

figure 1, which is made by the living beach worm *Paraonis fulgens*, was shown by Röder [36] to be a feeding burrow rather than a grazing pattern. The difference, in this case, is that the mucous that lines the burrow walls forms a 'trap' for the micro-organisms on which the animal feeds. After building the spiral (which is unfortunately washed away at each high tide) repeated visits to the lobes provide the worm with a continual harvest of food. The compact nature of the burrow may confer an advantage on the animal when the worm population is high [4]. Röder's analysis of the *Paraonis fulgens* meander as a feeding burrow inspired Seilacher [38] to review the functional interpretation of a number of trace fossils such as the honeycomb-like structures of genus *Paleodictyon*, an example of which is shown in figure 13. These traces break a 'cardinal rule' of foraging meanders: never recross the existing trail. With their regular patterning and multiple exits Seilacher suggests that these traces might be better understood as food "search nets" or as "farms" for bacteria or fungi. In [38] he also proposes a number of possible mechanisms that could be used to derive similar patterns, for instance, the superpositioning of two meandering traces at right angles to one another.

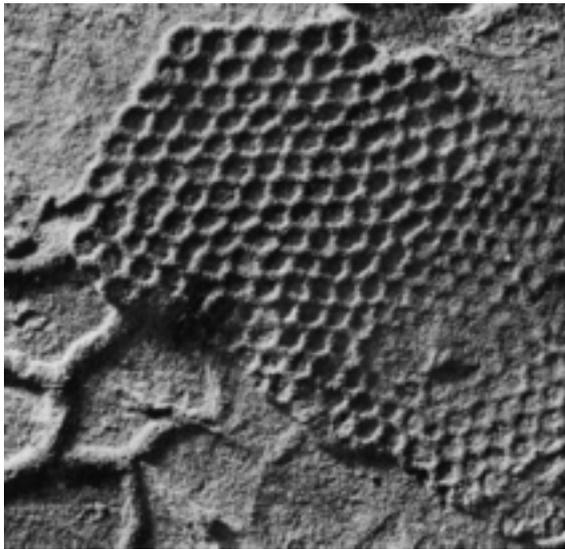


Figure 13. A honeycomb-like trace fossil of genus *Paleodictyon* (Upper Cretaceous period). Interpreted by Seilacher as a feeding burrow or "farm" for bacteria or fungi. Reprinted from [38] with permission of the Liverpool Geological Society.

Shortly after their publication, the trace fossil simulations of Raup and Seilacher attracted the interest of the mathematician John Conway (better known for inventing the game of 'Life') who with his colleague Michael Paterson, investigated an idealized 'worm' moving along the lines of an isometric grid. Their computer simulations,

described in [19], focused on classifying the trace patterns generated by worms moving on an infinite grid of equilateral triangles (i.e. six lines meeting at each node), some of which show a fascinating emergent structure.

The theoretical understanding of meandering behavior was also the focus of research by Papentin [32], who used a genetic algorithm to evolve the rules guiding movement of simulated worms on a rectangular grid of finite size. Each worm foraged for 140 steps and was evaluated by the number of steps which traversed new territory less the number which covered segments already traversed by the same or another worm. With high population densities (50 individuals in an area 100x100) Papentin found that, within 60-100 generations, spiraling and meandering trails evolved that were similar to those described for trace fossils. This work therefore supports the hypothesis that competition between individuals was a driving force behind the evolution of compact foraging patterns.

The guided meanders that we have been considering belong to the class of 'stigmergic' processes in which the behavior of the organism restructures its environment which in turn effects the activity of that organism (or other organisms) [42]. The phenomenon of stigmergy is most often explored in the context of group behavior such as the members of a termite colony working together on the construction of a nest. The regulation of behavior in such groups does not depend on direct communication between animals but is largely directed by the environmental structure which emerges from their 'collaborative' efforts. Unfortunately, there is little evidence of co-operative behavior in trace fossils, although in some cases, such as feeding burrows, it is possible that the structure is the work of more than one individual. Some trace-makers may not have distinguished their own trail from that made by another animal of the same species. For instance, Seilacher [38] describes a fossil which appears to show an animal whose tunneling behavior becomes entrained by the tunnel system of its neighbor. This form of stigmergy may perhaps have served as a pre-adaptation for the co-operative foraging and building activities that we see perfected in the social insects.

6 A 'Cambrian explosion' of behavior-based robots?

The experiments we have described show that a trace-making robot, controlled by a small number of reflexive behaviors, can generate a variety of fossil-like 'foraging' trails. More significantly perhaps, they demonstrate a striking similarity between the sensorimotor behavior of ancient, trace-making

animals, and that of simple, reactively-controlled, behavior-based robots. This similarity locates the behavior of such robots at a grade similar to animals of the late Precambrian or early Cambrianⁱ.

Deposits of *body* fossils from the early Cambrian show the first appearance of preservable, bilateral animals that can be definitively identified as members of the modern metazoan phyla. The base of the Cambrian, a period of perhaps less than twenty million years, appears to have seen an explosive development of many different body forms and complex nervous systems. The organisms of this fauna achieved a great diversity of methods of locomotion, had an abundance of different sensory mechanisms including compound eyes, and possessed a wide range of behavioral repertoires including predation [12, 30]. Some of these animals may therefore have lead very mobile and active lifestyles and exhibited complex and appropriate reactions to varied stimuli. Comparative and paleoneurobiological studies also indicate that 'groundplans' for the neural circuitry of the different phyla were established during this period, and that these placed significant constraints on subsequent evolution. For instance, the basic pattern of insect nervous systems was probably present in Arthropod ancestors of the Cambrian, and has since shown primarily quantitative rather than qualitative changeⁱⁱ. It also appears likely that the basic plan for the vertebrate nervous system was established at an early stage [8, 22, 25, 40], and probably within 100 million years of the initial Cambrian explosion. Miklos [30] who has referred to this period of rapid evolution as a "big bang" in the evolution of complex nervous systems, suggests that:

"Complex brains were unlikely to have been painstakingly 'wired-up' synapse by synapse over hundreds of millions of years. We are faced with the exciting prospect that nervous systems can be constructed rapidly". ([30] p. 854)

That the evolution of nervous systems in the early Cambrian metazoa proceeded at such a cracking pace should give encouragement to the designers of robot control systems. However, a number of cautions should be entertained with regard to the prospect of an imminent 'explosion' of behavior-based robots. First, we should recognize that robotics currently lacks building materials with the versatility and intelligence of the eukaryotic cell (itself the outcome of three billion years of evolution). Second, much work in the design of robot control circuitry is not far above the level of specifying individual synapses. To emulate nervous system evolution more closely, progress is needed in understanding and modeling the sophisticated development processes that control gene expression in neural circuitryⁱⁱⁱ. Finally, we might consider whether an explosion of behavior-

based robots has already occurred in that a wide variety of platforms that exhibit mobility and complex behavior have already been built and demonstrated (see [7, 29] for review), many of them going well beyond the 'wall-following' grade. This is not to suggest that we have achieved the morphological or neural complexity of the early Cambrian fauna, but we may be beginning to replicate their level of intelligent behavior.

Acknowledgements

The authors are grateful to Michael Port for technical assistance; to Graham Bull, John Hallam, Gillian Hayes, and Tim Colles for advice and help in building and debugging the robot; and to Leila Edwards for suggesting some improvements to the manuscript. Photographs of the robot trails were taken by Len Hetherington.

ⁱSeilacher [37, 39] has suggested that the complexity of trace fossil behavior increased gradually for 100 million years or so after the start of the Cambrian period; however, recent finds have caused this conclusion to be revised and it is now thought likely that diversity increased during the Cambrian radiation and has been relatively constant since [13, 33]. Certainly, efficient grazing patterns are now known from rocks of the Vendian and early Cambrian periods [13].

ⁱⁱEdwards [15, 16] describes the evolution of the nervous systems of insects as having been "astonishingly conservative", despite remarkable variations in body plans, with perhaps the most significant trend being towards miniaturization of the neural circuitry in some species, and the most variation being in the relative volume of sensory processing.

ⁱⁱⁱModelers and robot-builders are now beginning to take an interest in simulating the processes of neural development (see [26] for review). Progress in the area could have important consequences for the automated construction and evolution of complex control systems.

Appendix: The robot trace-maker

The robot trace-maker was assembled from a 'Lego vehicle' kit developed by the Intelligent Sensing and Control Laboratory in the Department of Artificial Intelligence at the University of Edinburgh, UK. The mechanical and structural elements of the robot are composed entirely from Lego Technic™ components, and are therefore easily reconfigurable. The onboard microprocessors and related electronics provide a programmable means of interfacing sensor and motor components. Control programs, written in the 'Control Process Language' [10] which simulates concurrent processing of behavioral modules, are compiled on an external computer and downloaded to the robot which operates autonomously thereafter. A description of earlier Lego vehicle technology is given in [14], a similar Lego robot to the one employed here has been customised by Webb to model the phonotaxis behavior of crickets and is described in [44, 45].

The overall robot dimensions are 40 cm long and 15 cm wide (excluding the LDR arms). The distance between drive wheels, which determines the turning radius, is 11 cm. The LDR sensors, on each side of the robot are positioned 6 cm and 14 cm away from the robot body, 2.3 cm above the floor, and 10 cm in front of the center of rotation (the center of the drive wheels). Rolls of 10 cm wide white tissue are used to generate the paper trail, the axle of the paper dispenser being positioned 15 cm behind the robot's center of rotation and 16 cm above the floor. The motor driving the dispenser operates at a fixed speed of rotation; however, the diameter of the paper roll decreases as the trail is laid, so paper is actually dispensed at a variable speed. During the advance and thigmotaxis behaviors paper is dispensed continuously at a speed close to the maximum forward velocity of the robot, when the robot moves more slowly the trail therefore takes on a layered appearance (see, for example, figure 4). No paper is dispensed during *phobotaxis* as pilot experiments indicated that this helps to prevent excessive layering of the paper trail. A start/stop button on the robot allows its progress to be instantaneously interrupted. This is necessary to replenish the paper supply, and, and on some occasions, to clear a paper 'jam' in the dispensing mechanism. In the experiments described these interruptions were infrequent, and were made at moments where they did not effect the subsequent trajectory of the robot.

Experimental sessions were held under conditions of ambient light (either from ceiling lights or from windows). Appropriate thresholds for thigmotaxis and phobo-taxis were determined at the start of each session by placing the robot so that one LDR sensor was above a sample of paper and another above the

(darker) carpeted floor. The threshold for thigmotaxis was set slightly below the reflectance of the paper and that for phobotaxis slightly above the reflectance of the floor. The LDR sensors used had a response range of 0-255. In a typical session, where the reflectance of the paper and floor were approximately 150 and 110 respectively, the thigmotaxis threshold was set at 140 and that for phobotaxis at 122.

References

1. Altman, J.S. and J. Klen, New models of motor control. *Neural Computation*, 1989. **1**: p. 173-183.
2. Beer, R.D., R.E. Ritzmann, and T. McKenna, *Biological Neural Networks in Invertebrate Neuroethology and Robotics*. 1993, San Diego, CA: Academic Press.
3. Braitenberg, V., *Vehicles: experiments in synthetic psychology*. 1986, Cambridge, MA: MIT Press.
4. Bromley, R.G., *Trace Fossils: Biology and Taphonomy*. 1990, London: Unwin Hyman.
5. Brooks, R.A., A robust layered control system for a mobile robot. *IEEE Journal on Robotics and Automation*, 1986. **RA-2**: p. 14-23.
6. Brooks, R.A., A robot that walks: emergent behaviour from a carefully evolved network. *Neural Computation*, 1989. **1(2)**: p. 253-262.
7. Brooks, R.A., New approaches to robotics. *Science*, 1991. **253**: p. 1227-1232.
8. Butler, A.B. and W. Hodos, *Comparative Vertebrate Neuroanatomy*. 1996, New York: Wiley-Liss.
9. Chiel, H.J. and R.D. Beer, The brain has a body: Adaptive behavior emerges from interactions of nervous system, body and environment. *Trends in Neurosciences*, 1997. **20**: p. 553-557.
10. Colles, T., J. Hallam, and T. Smithers, *Control Process Language Compiler User Manual*, 1994, Department of Artificial Intelligence, University of Edinburgh, UK., DAI Technical Paper No. 17a.
11. Connell, J.H., *Minimalist Mobile Robotics*. Perspectives in Artificial Intelligence. 1990, Boston: Academic Press.
12. Conway Morris, S., Burgess shale faunas and the Cambrian explosion. *Science*, 1989. **246**: p. 339-346.
13. Crimes, T.P., Changes in the trace fossil biota across the proterozoic phanerozoic boundary. *Journal of the Geological Society*, 1992. **149(Pt4)**: p. 637-646.
14. Donnett, J. and T. Smithers, Lego vehicles: a technology for studying intelligent systems, in *From Animals to Animats: Proceedings of the First International Conference Simulation of Adaptive Behaviour*, J.-A. Meyer and S. Wilson, Eds. 1991, MIT Press: Cambridge, MA. p. 540-549.
15. Edwards, J.S., One organism, several brains: evolution and development of the insect central nervous system. *Annals New York Academy of Sciences*, 1977. **299**: p. 59-72.
16. Edwards, J.S. and J. Palka, Insect neural evolution—a fugue or an opera. *Seminars in Neuroscience*, 1991. **3**: p. 391-398.
17. Franceschini, N., From insect vision to robot vision. *Philosophical Transactions of the Royal Society of London. B*, 1992. **337**: p. 283-294.
18. Frey, R.W. and A. Seilacher, Uniformity in marine invertebrate ichnology. *Lethaia*, 1980. **13**(183-207).
19. Gardner, M., Mathematical games: fantastic patterns traced by programmed "worms". *Scientific American*, 1973. **11**: p. 116-123.
20. Gaussier, P., Moving the Frontiers between Robotics and Biology. *Robotics and Autonomous Systems Special Issue*, 1995. **9(2-4)**.
21. Hallam, J.C.T. and C.A. Malcolm, Behaviour—perception, action, and intelligence—the view from situated robotics. *Philosophical Transactions of the Royal Society of London Series A*, 1994. **349**(1689): p. 29-42.
22. Hodos, W., Some perspectives on the evolution of intelligence and the brain, in *Animal Mind—Human Mind*, D.R. Griffin, Ed. 1982, Springer-Verlag: Berlin.
23. Hofmann, H.J. and I.M. Patel, Trace fossils from the type 'Etcheminian Series' (lower Cambrian Ratcliffe Brook formation), Saint John area, New Brunswick, Canada. *Geological Magazine*, 1989. **126(2)**: p. 139-157.
24. Huckriede, R., Eine spiralförmige Lebensspur aus dem Kulmkieselschiefer von Biedenkopf an der Lahn (Spirodesmos archimedeus n.sp.). *Palaontologische Zeitschrift*, 1952. **26**: p. 175-180.
25. Jerison, H., *Evolution of the Brain and Intelligence*. 1973, New York: Academic Press.
26. Kodjabachian, J. and J.-A. Meyer, Evolution and development of control architectures in animats. *Robotics and Autonomous Systems*, 1995. **16**: p. 161-182.
27. Maes, P., Behaviour-Based Artificial Intelligence, in *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, J.-A. Meyer, H.L. Roitblat, and S.W. Wilson, Eds. 1992, MIT Press: Cambridge, MA. p. 2-11.
28. Mataric, M.J., *A Distributed Model for Mobile Robot Environment-Learning and Navigation*, 1990, MIT AI Lab, AITR-1228.
29. Mataric, M.J., Behavior-based control: examples from navigation, learning, and group behavior. *Experimental and Theoretical Artificial Intelligence, special issue on Software Architectures for Physical Agents*, 1997. **9(2-3)**: p. 323-336.

30. Miklos, G.L.G., Molecules and cognition: the latterday lessons of levels, language, and Iac. *Journal of Neurobiology*, 1993. **24**(6): p. 842-890.
31. Nehmzow, U., T. Smithers, and B. McGonigle, Increasing behavioural repertoire in a mobile robot, in *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*, J.-A. Meyer, H.L. Roitblat, and S.W. Wilson, Eds. 1992, MIT Press: Cambridge, MA. p. 291-297.
32. Papentin, F., A Darwinian evolutionary system III: experiments on the evolution of feeding patterns. *Journal of Theoretical Biology*, 1973. **39**: p. 431-445.
33. Raff, R.A., *The Shape of Life: Genes, Development and the Evolution of Animal Form*. 1996, Chicago: Chicago University Press.
34. Raup, D.M. and A. Seilacher, Fossil foraging behaviour: computer simulation. *Science*, 1969. **166**: p. 994-995.
35. Richter, R., Psychische reaktionen fossiler tiere. *Palaeobiologica*, 1928. **1**: p. 225-244.
36. Roder, H., Gangsysteme von *Paranois fulgens* Levinsen 1883 (Polychaeta) in ökologischer, ethologischer und aktuopalaontologischer Sicht. *Senckenbergiana*, 1971. **3**: p. 3-15.
37. Seilacher, A., Fossil Behaviour. *Scientific American*, 1967. **217**: p. 72-80.
38. Seilacher, A., Pattern analysis of *Paleodictyon* and related trace fossils, in *Trace Fossils 2: Geological Journal Special Issue No. 9*, T.P. Crimes and J.C. Harper, Eds. 1977, Seel House Press: Liverpool. p. 289-334.
39. Seilacher, A., Evolution of behaviour as expressed in marine trace fossils, in *Evolution of Animal Behaviour*, M.H. Nitecki and J.A. Kitchell, Eds. 1986, Oxford University Press: Oxford.
40. Stahl, B.J., Early and recent primitive brain forms. *Annals of the New York Academy of Sciences*, 1977. **299**: p. 87-96.
41. Steels, L., The artificial life roots of artificial intelligence, in *Artificial Life: an Overview*, C.G. Langton, Ed. 1995, MIT Press: Cambridge, MA. p. 75-110.
42. Theraulaz, G. and E. Bonabeau, Modeling the collective building of complex architectures in social insects with lattice swarms. *Journal Of Theoretical Biology*, 1995. **177**(4): p. 381-400.
43. Valentine, J.W., Late precambrian bilaterians - grades and clades. *Proceedings Of the National Academy Of Sciences Of the United States Of America*, 1994. **91**(15): p. 6751-6757.
44. Webb, B. *Modelling biological behavior or 'dumb animals and stupid robots'*. in *Second European Conference on Artificial Life*. 1993.
45. Webb, B., Using robots to model animals: a cricket test. *Robotics and autonomous systems*, 1995. **16**: p. 117-134.