Braitenberg Simulations as Vehicles of Evolution

Correy Allen Kowall
Department of Computer Science
Northern Michigan University
1401 Presque Isle Ave.
Marquette, Michigan 49855
(906)360-4750

ckowall@nmu.edu

ABSTRACT

In this paper, we describe a series of simulations that serve as a verification of the abstract similarity between vehicular and animal navigation. Valentino Braitenberg used this similarity to illustrate that vehicles controlled by very simple biologically inspired circuits, manifest a wonderful diversity of complex animal behaviors. By constructing a series of experiments that are designed around the possibility of interchanging phenomena that affect vehicles or animals, we hope to show that this analogical similarity is a useful tool.

Categories and Subject Descriptors

I.2.9& I.6.8 [**Computing Methodologies**]: Simulated Evolution of Vehicles – *animation, distributed, parallel, autonomous vehicles.*

General Terms

Algorithms, Experimentation.

Keywords

Evolution, divergence, cooperative populaces, adaptive systems, interference of adaptive systems.

1. INTRODUCTION

Numerous navigation problems have served as testing grounds for adaptive systems in the existing literature and software base [2][4]. We utilize the description of theoretical terrestrial vehicles offered by Valentino Braitenberg in his book *Vehicles: Experiments in Synthetic Psychology*[1]. In addition to the vehicular model we will utilize two salient points from this work: simple circuits generate complex behaviors and navigation by animals and vehicles is similar. It is perhaps important to note that these were not Braitenberg's central point. In fact it is clear his intent was to contrast the simplicity of designing a solution, with the difficulty of identifying a specific mechanism that a biological nervous system implements to solve the same problem. To illustrate this point Braitenberg based each of the chapters in the book on a theoretical vehicle with a simple biologically inspired control circuit that would manifest a familiar but complex animal behavior. If the analogical similarity of vehicle control and animal behavior is valid, a simulation that contains this type of vehicle model should be able to facilitate scenarios where biological phenomena emerge in vehicles or vehicular phenomena emerge under biological conditions.

In the first experiment we will attempt to create a scenario where a group of vehicles demonstrates the same tendency toward self propagating behaviors that many natural populaces exhibit. In the second experiment we will attempt to simulate the interaction between two biological phenomena by applying both to vehicles in a simple environment. Specifically the phenomena we simulate are the improvement of an individual by a learning algorithm and the improvement of a populace by environmental pressures. A secondary purpose of this experiment is to demonstrate the degree to which a simulation that is otherwise a simplification (wheels instead of legs etc.) can host complex interactions. The third experimental simulation was designed to be a reprisal of the environment in the second experiment with the addition of a vertical dimension. Fortuitously, when the simulation was allowed to run for several hours the populace demonstrated both a biological and vehicular phenomena: Speciation and Trains.

1.1 Vehicle of Choice

Braitenberg vehicles of Type 2 (Figure 1) are described as terrestrial vehicles (operating on a plane) that have two wheels which are independently controllable producing the net effect of directional and velocity control of the entire vehicle. The cars can
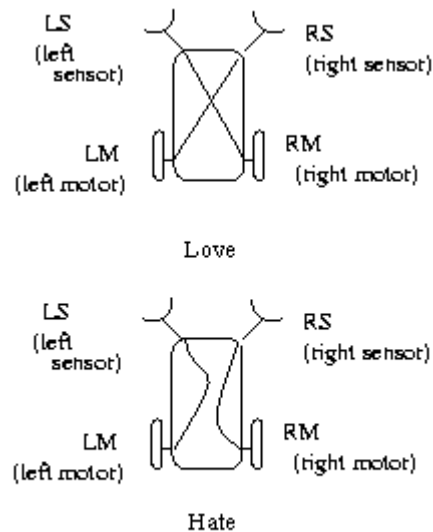

Figure 1. Characterized Vehicles.

be mathematically modeled as system of two functions (one for each wheel) on some combination of the variables that represent

environmental information. In these models directional change is equivalent to the difference of the values produced by the pair of functions, and the sum of the values represents travel or magnitude. Cars of Type 3c are similar to Type 2 but have any number of specialized sensors to resolve the environment they inhabit. With a dedicated simple sensor it is easy to simplify or circumvent the process of cognition in our models.

1.2 The Paths to Improvement

Learning in artificial neural networks and simulated evolution are both examples of improving systems. Computerized vehicle simulation gives us the capability to implement both of these techniques. At the level of an individual vehicle a neural network can be 'trained' or changed in a way that produces improvements of individual behaviors. At the population level, evolutionary operators like transmission of characteristics, mutation, or pressure exerted by environmental conditions can result in improvements of the performance of a navigational task. Performance in individuals and whole populations can be measured by the same environmental variables like linear distance / time, etc. Furthermore with the utility .of a flexible set of reliable navigational testbeds, experimentation with improving systems can be benchmarked across disparate categories of navigational scenarios[4]. Computerized simulation allows complete transparency and control of the set of information that constitutes an artificial environment. Furthermore an experimenter is at liberty to manipulate access to that information by instantiated objects operating in the simulated environment. These two controls allow us to benchmark the number of exposures to stimuli an improving system receives before performing a behavior. Additionally graphical rendering of the computerized simulation offers the possibility of repeated visual demonstration of emergent behaviors or morphological changes. Vehicles are guided by independent instantiated objects created from behavioral classes which serve to compartmentalize variable information in both simulation engines. This fact not only facilitates processing collision and sensation determinations, it guarantees uniformity of test conditions.

1.3 Methods

Two simulators were employed during these experiments: A Java based two dimensional vehicle simulator created by Shawn O'Neil, and Conitec's 3d GameStudio version A6. The GameStudio itself is proprietary but an executable version of the 3d simulation, and an appletized version 2d simulator are available at http://myweb.nmu.edu/~ckowall. In all comparative or combined simulations vehicle properties were held constant in that the types of sensors and means of expressing affect on the environment were identical.

2. EXPERIMENTS AND DEMONSTRATIONS

2.1 Contagion of Behavior.

Memes and virii have similar transmission operations in that they require a host, as well as temporal and spatial proximity to propagate. Braitenberg characterized his vehicles' behaviors with anthropomorphic titles like 'Love' or 'Hate' to describe 'seeking' or 'avoiding'. This lead to the conjecture (on the part of the investigator) that the most successful, or likely to proliferate, behavior modifying virulent effect is Braitenberg's 'Love'. Because a host infected with the characteristic of 'Love' will tend to transport and transmit the effect to others that can serve as a host, it will be particularly virulent. To test this we created a

simulation by modifying the Java based Braitenberg vehicle simulator. In this scenario all cars sense and carry blue lights; the only environment cues consists of blue lights that are carried by other vehicles. If two vehicles are at the same location at the same time the one that has traveled further will transmit its look up table (the cars are controlled by a discretized mechanism) (Table 1) to the other vehicle. Physical morphology was held

Table 1. Look up table for clockwise "Orbiter" vehicle

Left sensor	Right sensor		Action
0	0		Turn right
0	1		Go forward
1	0		Turn left
1	1		Turn left

constant and the two conic sections that represented the area of sensation (of blue lights) had an area of mutual inclusion (analogous to predators' eye positions). This overlap is important because it allows the deduction of additional location data by the control mechanism. If an example of all 256 possible vehicle look up tables that can represent behavior driven by two binary stimuli (eyes) acting on two binary responses (wheels) are placed into the same environment the population becomes dominated by a broad class of 'seekers'. 'Orbiter' and 'Wall-follower' behaviors came in a distant second but persisted well into many trials because the cars were mobile enough to avoid seekers and accumulate travel credit. Other persistent population elements included 'Travelers' or 'Traveling Avoiders' that move forward without stimuli and turn away from those lights they do encounter. One could easily speculate that if the width and offset of the cones of sensation could be transmitted or mutated that speciation along the lines of predators and prey would occur. In the third experiment we will show just that type of divergence.

2.2 Interference by Improvement Systems.

Subsequent efforts included developing code that allowed cars to be operated by feed forward multi-layer neural networks. Training algorithms including back-propagation of errors method, an on-line adaptive resonant training algorithm, and natural selection (randomized static networks selected by performance and environmental conditions for representation in future populaces) were all tested [3]. All neural cars had four input, four hidden, and two output nodes with hyperbolic tangent threshold functions. Because the operation of selecting members of the populace for replication does not interfere directly with the on-line selection of connection weights in the network that guides individual vehicles it was possible to test the effect of their operation on a single populace over the same period of time. The navigation task was composed of: starting at one of a hundred starting points, avoiding a barrier composed of red lights (that would destroy vehicles that collided with them), ending at a set of blue lights (that count cars that collided with them). All cars were configured with two sensors for red lights and two sensors for blue lights. First we tested a set of 100 look up table driven cars of the same sensor layout as the neural model of which 0 succeeded. This result is easily analyzed, because the pan and breadth of the cone of sensitivity of the red (avoidance) sensors spanned the narrow opening in the barrier composed of red lights. The discrete control architecture is not capable of temporarily 'ignoring' the negative stimuli on the way to the goal so vehicles with this type of 'brain' usually come to rest jammed in the opening refusing to advance but remaining fixed on the goal.

Second we tested a set of 100 neural network controlled cars with random connection strengths and analog sensors. On average three of 100 random weighted networks guide their vehicles to the blue lights. Third we tested an on-line trained network identical at incept to the cars with randomized starting weights. Feedback was provided to the training algorithm by signal from the sensors themselves. If a behavioral change coincided with an increase of signal from the 'goal' sensor, it is reflected in a reduction of the size future weight changes. Conversely the increase of signal from 'obstacle' lights results in an increase in the step size of future weight changes. On average seventeen of 100 cars with on-line learning reached blue lights within 20000 simulation cycles. In the fourth test 100 untrained randomized neural cars were used to isolate an average of three successful weight configurations that were proportionately represented in a populace of 100 offspring that were then started at each of the 100 starting points yielding an average of 10 successful navigations. In the fifth test both selection and on-line training were applied to the same cars. After one generation of elite selection 15 of 100 trained cars completed the navigation task. Although the results are inconclusive, they essentially show the two methods of improving the populace interfered with each others' operation in spite of the fact they do not directly interact. It is important to recognize that the initial connection weight state of the second or selected generation was not random but of the *starting states* of successful members of the first populace. In the sixth test 100 initially randomized connection weight configurations were tested and selected producing an average of three cars that were then proportionally represented by 100 offspring which were trained by the online learning while being tested. Strangely the second populace failed to reach the success rate of the untrained selected populace, by producing only an average of 7 successful cars. It is tempting to say that effects like overspecialization by selection (combined with of the variance in starting position) and the non-linear effect of the on-line learning algorithm produced the observed loss of improvement performance when improvement algorithms are combined (test 6 & 7) but a causal explanation is beyond the scope of this paper. Future work will include tests of Lamarckian style transmission of trained weight strength configurations that may offset the observed interference.

2.3 Emergent Specialization in Small Populaces.

In the third experiment, a navigational scenario, like the one employed in the second experiment (complete with small opening) was created with Conitec's 3d GameStudio World Editor. In this test the wall separating the starting position of the cars and the goal locations did not destroy the car but instead operated as a simple obstacle. For the sake of simplicity cars were controlled by a look up table. The genome of the three dimensional cars includes physical features like sensor angle and range as well as values in the output column of the look up table (the rotational velocity of a wheel given a sensor state). An aggregate fitness function was composed of functions on variables like the average distance to the objective and a constant valued reward for reaching the goal. Selection was based on randomly selected tourney style elimination and replication. Both the elimination and replication were executed when the population was in a target range from 25-50. Because Conitec's SDK allowed rapid modification and observation by means of sending a virtual avatar to 'witness' the scene of the cars operation, many different modifications to the fitness function,

evolutionary algorithm, and navigation task were tested. In one long running scenario two characteristic subgroups emerged: wall/traffic jam avoiders with widely separated collision sensors, and followers that had crossed collision detectors with reversed logic such that when they were born (at the same location) following an avoider they would lock their sensors on the back of that vehicle and follow it to the goal! (Figure 2) It appears that we have observed the emergence of train cars amongst a populace of cars. Although the phenomena is unstable and suffers severe periods of negligible performance (in terms of the populaces ability to reach the goal) it is perhaps of more interest to observe emergence than improvement.

Figure 2. Emergent Cooperative Species.

2.4 Future Work

Future work will concentrate on using locally (node) referenced records of neural activity as an input set for networks that adjust the connection weights in an underlying behavior response neural network, and using evolutionary scenarios to select amongst populations of vehicles with meta-level networks.

3. ACKNOWLEDGMENTS

4. REFERENCES

[1] Braitenberg, V. *Vehicles: Experiments in Synthetic Psychology.* The MIT Press, Cambridge, MA,1986.

[2] Floreano, D. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines, A* Bradford Book, Cambridge, MA, 2004

[3] Haykin, S. *Neural Networks: A Comprehensive Foundation.* Prentice-Hall, Upper Saddle River, NJ, 1999.

[4] Kennedy, J. and Eberhart, R. *Swarm Intelligence.* Academic Press, San Diego, CA, 2001.

[5] Spector, L. and Luke, S. *Cultural Transmission of Information in Genetic Programming.* Proceedings of the Cognitive Science Society Meeting, 1996.