Defining Emergent Behavior for Mobile Robots

Daniel P. Stormont

Department of Electrical and Computer Engineering, Utah State University, Logan, Utah

Abstract

One of the claims made for behavior-based mobile robots is that they are able to cope with environmental uncertainty more robustly than planning-based robots because of emergent behaviors. The problem is that there are almost as many definitions of what characterizes emergent behaviors as there are researchers investigating behavioral robots. If there is to be any chance for predicting or characterizing emergent behavior, the starting point should be a definition that can be applied to the domain of behavioral mobile robots.

This paper proposes a definition for emergent behavior for mobile robots. It begins by considering what emergent behavior might look like, followed by historical claims of emergent behavior in mobile robots. Then some examples are examined to try to clarify the aspects of emergent behavior. Finally, a definition for emergent behavior is proposed which can be used as a starting point for future research, as discussed in the last section.

Keywords: Emergence, behaviors, robot, behavioral robot, mobile robot.

1. Introduction

An often cited advantage of behavioral autonomous robots and robotic swarms is "emergent behavior". Nearly every author who has written about behavioral robots has used this term. Unfortunately, there are nearly as many definitions for emergent behavior as there are authors writing about it.

Why is a definition of emergent behavior important? The simple answer is, if emergent behaviors in behavioral robots are ever going to be understood and anticipated or predicted, it is essential to have a working definition of emergent behavior as it applies to behavioral robots. Predicting emergent behavior is desirable when designing a behavioral robot in order to have a complete understanding of the resulting design. Predicting emergent behavior is absolutely necessary for critical applications of autonomous robotics. For example, as unmanned combat vehicles become more prevalent and more autonomous, it is essential that any behavioral components added to the robots to enable them to deal with complex and unpredictable environmental stimuli be thoroughly understood and any possible emergent behaviors anticipated. [1] It doesn't take much imagination to understand why an unanticipated behavior from an armed robot would be undesirable.

The Solarbotics Photopopper robot shown in Fig. 1 is an example of a behavioral robot that may be considered to display emergent behavior. [2] The robot has two light sensitive eyes that allow it to follow a light source. If the light source is stationary, the robot will orbit the light source such that the focus of its orbit is at the brightest point of the light source. If the robot encounters an obstacle keeping it from the light source, as detected by two feelers on the front of the robot, it will slowly work its way around the obstacle by alternately turning away from the obstacle and turning toward the light source. Surely if one observed an insect exhibiting these behaviors, it would be tempting to call them purposeful, even intelligent, in a simplistic, but determined, manner.

Fig. 1: A Solarbotics Photopopper orbiting a light source.

However, if the circuit for the Photopopper is analyzed (see Fig. 2), it becomes apparent that the apparently purposeful motion of the robot can be explained as the merging of a number of different sensorimotor responses in the motor control circuitry of the robot. The motor opposite one of the photosensitive "eyes" will receive a driving current proportional to the strength of the light detected by the eye. Likewise, the motor opposite to one of the feelers that has been closed by coming into contact with an obstacle will turn off, causing the robot to turn away from the obstacle. The interaction between these separate behaviors and the environment results in the global behavior observed from the robot.

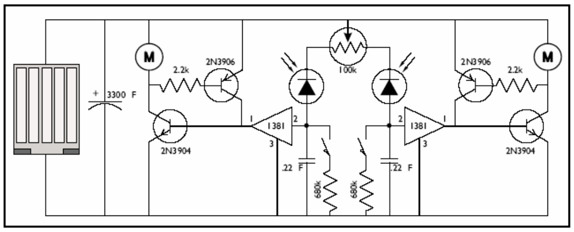

Fig. 2: The Solarbotics Photopopper control circuit. [2]

So, if the behavior can be easily understood, even if the global behavior was not explicitly designed in, is it an emergent behavior? That is the question this paper will attempt to answer.

2. Claims for Emergent Behavior

There have been many claims of observed emergent behavior in autonomous agents. The earliest claim for emergent behavior in a robot appears to be in [3]. In this book, Walter described the development and testing of robotic tortoises, which he gave the whimsical name of *Machina speculatrix*. Two of these robots were eventually built and tested in a closed environment. A restored version of one of these robots appears in Fig. 3. These robots were observed to have some truly remarkable behaviors: they would explore the environment in search of a light source that was sufficiently bright to be of interest, but not so bright as to repel the robot, they could navigate around obstacles and proceed toward the light source, and they could return home to recharge when they became "hungry". The remarkable thing about these robots was that, like the Photopopper discussed earlier, they didn't have any of these behaviors explicitly designed in. In fact, their control circuitry was a model of simplicity itself. They were controlled by two vacuum tubes and a light sensor that drove a steering motor and a propulsion motor. All of the observed behaviors were a result of the biasing of this control circuit in response to environmental stimuli.

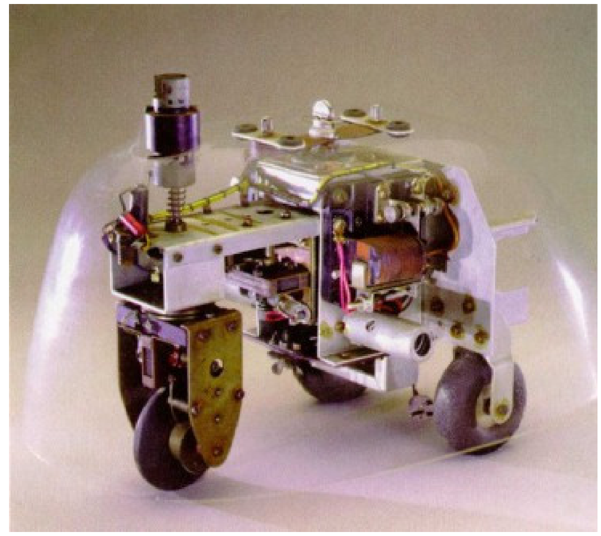

Fig. 3. A photograph of *M. speculatrix* restored to its original condition. [4]

Nearly a quarter of a century later, Braitenberg explored some concepts of neuroscience using a series of hypothetical vehicles in [5]. These vehicles, as Walter's *M. speculatrix* had done experimentally, provided examples of complex observed behaviors arising from very simple rules. In fact, Braitenberg went so far as to attribute emotional motivations to the observed behaviors of many of the vehicles he described. It seems likely that many observers of these vehicles would share Braitenberg's attribution of motivation upon observing the robots operating in their environment. The simplest vehicle described in [5] was termed a Type 1 vehicle. It consisted of nothing but a single motor connected to a single light sensor. The speed of the motor would be proportional to the amount of light received by the sensor. Even without a steering mechanism, if the robot was small enough to be affected by environmental factors, such as friction and grade, it would appear to move erratically with Brownian motion. An observer might be tempted to call it restless and, if it were properly disguised, would almost certainly call it alive – if not particularly intelligent. With the addition of additional motors and sensors, the vehicles demonstrated other emotions, such as fear and aggression (avoiding or attacking a light source), love (orbiting a light source), and a series of more complex emotions as threshold circuits, memory, and reinforcement learning were added. While some of the vehicles would have required circuitry that doesn't currently exist, many of them could be implemented.

Brooks, Braitenberg's contemporary at the Massachusetts Institute of Technology, was thinking about ways to utilize emergent behaviors resulting from simple rules to develop a new methodology for building robots. His subsumption architecture was first presented in [6]. The subsumption architecture used the concept of independent sensorimotor behaviors (which Brooks called competencies) interacting with each other through an arbitration mechanism. In this way, those behaviors most essential to the robot's safety and survival could take over when needed, allowing higher-level reasoning to occur at other times. Since the behaviors all acted independently, higher-level functions could continue to be added to the robot without the need to modify the existing competencies. An example of a subsumption architecture implementation is shown in Fig. 4.

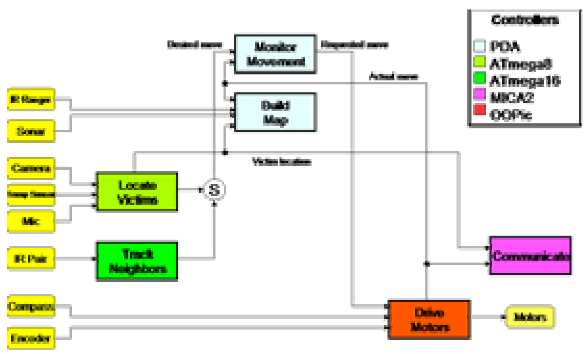

Fig. 4. The subsumption architecture diagram for the software on Utah State University's (USU's) Blue Swarm 3 search and rescue robots.

Finally, Tilden, after hearing a lecture about the subsumption architecture by Brooks, wondered whether it would be possible to implement the subsumption architecture entirely with analog circuitry. The resulting nervous-net architecture would become the basis for BEAM (Biology, Electronics, Aesthetics, and Mechanics) robots, as described in [7]. The Solarbotics Photopopper presented in the introduction is an example of a BEAM robot, where the emergent behavior is a result of varying biases in the control circuitry resulting from the robot's interaction with its environment.

3. Examples of Emergent Behaviors?

The historical examples from the last section were all examples of behaviors emerging from a single robot as it interacts with its environment. Emergent behavior can be even more difficult to predict and to identify when a number of behavioral robots interact with each other. An example of this is shown in Fig. 5, which shows a simulation of a swarm of six search and rescue robots built at USU, each of which are subsumption architecture behavioral robots in their own right, that interact with each other to explore a simulated disaster environment. The resulting behavior, though none of the robots were programmed to explicitly explore the environment, was that the swarm covered all of the accessible areas of the rescue arena consistently during the American Association for Artificial Intelligence (AAAI) Rescue Robot competition in 2001.

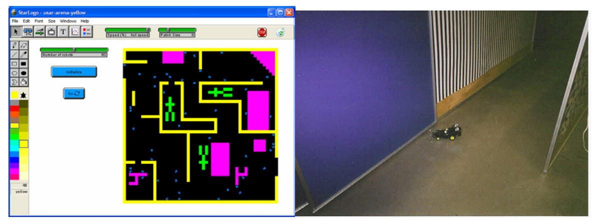

Fig. 5. A simulation of Blue Swarm 2 robots and a photo of one of the swarm robots exploring the rescue arena at the AAAI 2001 Mobile Robot Competition.

The common thread for all of the examples presented so far, is that they all exhibit behaviors that were not explicitly designed into the robot. From simple controls come complex actions. So, are these emergent behaviors?

In [8], the argument is made that emergent behavior can be defined by the "surprise" factor. If you design and implement a robot and then it does something that is surprising, then that surprising behavior is an emergent behavior. A quantitative value for the percent of emergence could be produced by gauging the degree of surprise at the unexpected behavior. The problem with this definition, as pointed out in [9], is that the element of surprise is subjective. An observer with no robotics knowledge could observe an animatronic robot at an amusement park and be surprised by its behavior – even if the behaviors are completely scripted and repetitive. If it is possible for an emergent robot to become non-emergent if the observer learns how it works, then it seems doubtful that this is a useful definition.

A better, but perhaps too limited definition of emergent behavior is provided in [10]. It states that emergent behavior can be defined as "the global behavior of the robot or organism as a consequence of the interaction of the active individual behaviors." There seems to be little doubt that the interaction of behaviors is a critical component of emergent behavior. However, this definition could be applied to a robot that does not display any behaviors beyond the programming of its individual behaviors. For example, a robotic arm programmed to execute a pick behavior, followed by a place behavior that is executed after a signal (interaction) from the pick behavior that it has gathered a part would seem to fall into this definition, but is the robot displaying emergent behavior?

The definition of emergence in [9] contains another piece of the puzzle. It states that emergent properties are those recurring patterns that are regularly associated with events of interest. This pattern of producing behaviors in response to internal and external events would certainly be anticipated as part of emergent behavior, and necessary if there is going to be any possibility of predicting emergent behaviors. But this definition doesn't include the concept of simple rules producing complex behaviors, although this concept is discussed elsewhere in [9].

4. A Proposed Definition of Emergent Behavior

So, given the necessity of predicting emergent behaviors and the desirability of anticipating and understanding emergent behaviors, what could be used as a working definition of emergent behavior in robots? For the purposes of additional research into this area, I propose the following:

Definition. *Emergent behavior is an observed action or sequence of actions, not explicitly designed into the robot, resulting from the internal interaction of the robot's behaviors and their external interaction with the robot's environment.*

How could this definition be applied? Let us return to the Photopopper robot described in the introduction. Recall that the Photopopper had two behaviors built into its analog circuitry: one that would supply current to the motor on the opposite side of the robot from one of the light sensors proportional to the amount of light received by the sensor and one that would turn off the motor on the opposite side of the robot from a touch sensor switch that had closed. There is no explicit behavior for seeking out and orbiting a light source, nor is there a behavior for moving around an obstacle while alternately attempting to move toward a light source. Yet this is exactly what is observed when the robot is placed in an environment containing a light source and one or more obstacles. Applying the definition of emergent behavior, it is apparent that the interaction of the internal behaviors and environmental stimuli has resulted in a sequence of actions that were not explicitly designed into the robot. Thus, the definition above validates the assertion made in the introduction that the Photopopper is exhibiting emergent behavior.

5. Future Work

At this point, a valid criticism could be made that undesirable behaviors, like programming bugs, could be defined as emergent behaviors. And indeed, they are. The next step in this research is to use the proposed definition to develop a methodology for identifying, and more importantly, predicting emergent behaviors – both desirable and undesirable. The first step will develop a methodology for a behavior-based quadrupedal walking robot to predict emergent gait behaviors in rough terrain. The follow-on research will attempt to extend the methodology to predict the emergent global behavior of a swarm of interacting, behavioral search and rescue robots.

6. Acknowledgment

The author would like to thank the Space Dynamics Laboratory at Utah State University for their support with the Tomorrow PhD Fellowship.

7. References

[1] J. McLean, "AI & IA: The Use of Artificial Intelligence in High Assurance Systems", presented at the Fourteenth Conference on Artificial Neural Networks In Engineering (ANNIE), St Louis, MO, 2004.

[2] http://www.solarbotics.com.

[3] W. G. Walter, *The Living Brain.* London, England: Gerald Duckworth & Co. Ltd., 1958.

[4] http://stcroixstudios.com/wilder/fastkarl/drgreywalter.html

[5] V. Braitenberg, *Vehicles: Experiments in Synthetic Psychology*. Cambridge, MA: The MIT Press, 1984.

[6] R. Brooks, "A Robust Layered Control System for a Mobile Robot", *IEEE Journal of Robotics and Automation*, vol. RA-2, pp. 14-23, April, 1986.

[7] B. Hasslacher and M. Tilden, "Living Machines", Los Alamos National Laboratory, Los Alamos, NM, Report #LA-UR-94-2636, 1994.

[8] E. Ronald, M. Sipper, and M. Capcarrère, "Design, Observation, Surprise! A Test of Emergence", *Artificial Life*, vol. 5, no. 3, pp. 225-239, 1999.

[9] J. Holland, *Emergence: From Chaos to Order*. Reading, MA: Addison-Wesley Publishing Co., 1998, pp. 42-45.

[10] R. Arkin, *Behavior-Based Robotics*. Cambridge, MA: The MIT Press, 1998, p. 24.