Evolving Interaction in Robots

Andrew Sheng

The web of life forms a symphony of interaction and communication. Consider all the interactions a creature will perform and undergo throughout its life. A simple glance around one's surrounding yields many examples of such activities – a human speaks a string of words, a bee performs a complicated dance, a dog urinates onto a hapless bystander's leg. All of these actions carry an intention to convey some sort of message to others – EXAM TERRIFYING, FOOD THERE, MY TERRITORY. Entire academic fields, such as sociology, dedicate themselves to analyzing these behaviors.

The means by which such interactions may have arisen during the emergence and evolution of life remain one of the mysteries of the modern world. This paucity of knowledge arises from the problem that experiments regarding emergent behavior are extraordinarily difficult to conduct due to the lack of simple sample organisms for experiments. Case studies with individual microorganisms would be difficult to analyze and control. On the other hand, higher organisms tend to have generative cycles on the magnitude of months, years, decades – experiments with thousands of generations would require an extremely dedicated multigenerational research team. Furthermore, the distant past offers little helpful information – social structures are not easily fossilized. Given the lack of a suitable sample population, one faces the question: Is the quest for the origins of social interaction a limited one restricted to the analysis of the behavior of contemporary lifeforms? Does the lack of a suitable medium render the simulation of communication evolution an impossibility?

Unfortunately, the natural world does not seem to yield any optimal organisms. There exist few creatures quick to reproduce, simple to observe, easy to mutate, and, perhaps most importantly, free of preexisting behaviors hardcoded by genetics [1]. But, on a brighter note, when one considers models from fields other than biology, there does indeed exist

> **"The robots could inform their companions of food and poison"**

such an "organism" - a software algorithm! An algorithm is merely a set of easily replicable instructions; it also lacks the evolutionary baggage of billions of years of history. When placed into the body of a robot, the algorithm may play out the behavior coded by its instructions.

Although the replication of digital information is trivial (much to the despair of anti-piracy groups), the simulation of the process of biological evolution is a much more difficult matter. Unlike a cell, which can often survive a malformed protein or damaged DNA, a computer program can be easily destroyed by the introduction of even a single error. As a result, the mutable gene must not be the program itself, but instead some outside information that is interpreted by the program.

One implementation of this concept is a model called a "neural network." In a neural network, a program makes use of an amount of premade data as a blueprint (or gene) for the construction of a field of simulated neurons. Each simulated neuron is governed by two factors: connections to other neurons and some rudimentary calculative ability (such as determining whether the sum of a set of numbers is above or below some given threshold). The network is "run" by inputting data into some neurons. These neurons proceed to process the data and then output it to their neighbors, each "layer" of neurons making use of their internal instructions to interpret the data. To gather output, data is extracted from the output of several select neurons. Since the topology and properties of the network are entirely determined by its "genes", which exist independent of the host program, a neural network may be transmogrified via modifications to its blueprint. This mechanism allows the simulation of natural selection (and by extension, evolution) via the random and/or selective mutation and breeding of neural networks [2]. In sum, a robot serves as a simulation of a primitive organism. One such population of organisms resides in Switzerland, where flashing blue robots fight over glowing red floor tiles [1,3].

In an attempt to simulate the emergence of communication, researchers at the Ecole Polytechnique Federale de Lausanne have constructed a field with two small zones – one marked as "food," the other marked as "poison," both identical unless observed from a close distance. A population of small robots equipped with cameras and lights formed the denizens of this field. Each robot was guided by two factors: internal neural networks and an overarching rule of "food good, poison bad" [1].

The research began with the robots randomly flashing their lights – none of them could "understand" the speech of any other, only the local presence of food and poison. They were then left to wander randomly, eventually discovering food by the sheer mechanism of trial and error. The simulation would finish after a preset amount of time. Afterwards, the researchers ranked the robots by their success at collecting food and avoiding poison. The "genes" of those robots best at collecting food were then mixed together (in an approximation of mating) and randomly mutated. The new genomes were then replaced into robots in order to continue the simulation further.

More than five hundred "generations" of collection and mutation later, the robots (more specifically, the neural networks controlling them) were found to have developed the ability to utilize their cameras and lights so that they could inform their companions of food and poison. Some populations of robots evolved the tactic of flashing their lights when near food (as an invitation), while others leaned toward the tactic of flashing lights when near poison (as a warning). In other words, a group of robots had autonomously developed the ability to communicate – a development


The First Step to a Robot Society? Reproduced from [5]

made even more curious by the fact that earlier, the robots had been largely unaware of the existence of their fellows [1]. Thus was born the behavior of cooperation in order to take advantage of the strength of numbers – a single robot working alone cannot locate "food" as fast as a group of machines working together as a team.

However, the end result is not that of a robotic utopia. One aspect of the playing field was that the food "zone" was too small to support the entire robotic population; the robots would be required to push away others in order to acquire points toward their own placement in the next generation. As a result, the robots did not form a very harmonious society. Instead, some robots would send misleading messages to others in a "selfish" attempt to fool others in order to lure them away from the food. For example, in a population of robots with the "lights-mean-food" protocol, rogues would develop the tendency to cast signals over barren ground or poison in an attempt to fool others in order to increase their own chances of a free food zone. The mutation-based appearance of such rogues would often destroy cooperative groups of robots, as "survivors" would generally evolve to become less inclined to trust the signals of other robots [1].

A later study repeated under similar circumstance with similar robots found the existence of another deceptive behavior – that of withholding useful information from others. In many test "cases," most robots would often move to the strategy of avoiding using their lights while in the process of collecting food – this would prevent others from noticing anything special about that particular patch of land. However, the researchers found that in no case did the robots completely cease all usage of their lights; instead, even the most xenophobic machines would make some use of light, perhaps due to the fact that the marginal reward for "take-but-don't-give" is lessened when all robots refuse to share useful information [3]. Therefore, even though the robots fell to selfishly firing each other lies of omission, the robotic society did not completely degenerate into a purely competitive environment.

The study is admittedly simplistic. After all, human societies do not quickly degenerate into anarchies upon the appearance of criminals and con artists – there exist mechanisms to punish humans who scream "that's food" while pointing at cyanide. Furthermore, it is possible that the experiment unfairly promoted the benefits of pure competition over cooperation; the robots could be able to discover a more cooperative strategy in a different environment [4].

The Lausanne study reveals some interesting insights. Social behaviors commonly associated with living organisms are not restricted to life. Instead, it is likely that many behaviors are simply evolutionary responses to various environmental pressures (i.e. a given behavior could have developed as a random, yet beneficial, trait-guided action). Such an action could have thus improved an organism's fitness enough to pass it on the succeeding generations. In a frontier world, cooperation would help one's companions gather resources, while in a civilized place, deception could allow one to gain at the expense of others [1]. The fact that blocks of silicon and metal can spontaneously develop the ability to cooperate and cheat helps to bring out the more human question of exactly how many of one's actions are consciously generated, and how many are solely due to evolutionary psychology.

The emergence of the various robotic behaviors suggests that early societies (that is, of microorganisms, probably not humans) could have been highly dynamic, rapidly shifting between cooperative and competitive behaviors (depending on which one would be more useful in a given situation) before settling at some sort of equilibrium [4]. A resident of the primordial ooze would probably have been in a similar situation as the robots; it would have just had a simple predefined set of behaviors, anything else would have been developed later.

The results of the study could also provide guidance to those seeking to develop complex systems composed of many independent actors – whether these actors are robots or something else. The sheer variety of different behaviors developed by the robots during their evolution indicates that seemingly simple, decentralized systems may give rise to very complex behavior. Whether this complex behavior is desirable probably varies on the situation at hand.

The Lausanne study is just one case in which robotics is being put to use in an unconventional situation to advance the development of the understanding of the development of interactive behavior. Although there still exist significant differences between organisms simulated on silicon and organisms in flesh, it is conceivable that one day such experiments could guide humanity to an understanding of that composition of that symphony of interaction and communication, the web of life. ■

Andrew Sheng is a computer science major in CMU's class of 2013. He is most likely a member of the species Homo sapiens, however,classifying organisms by morphological features is not always a reliable method of identifying species.

References

1. Floreano, Mitri, Magnenat, Keller. Evolutionary Conditions for the Emergence of Communication in Robots. Curr Biol. 2007; 17:514-519.
2. Nolfi, Parisi. Evolution of Artificial Neural Networks. In: Arbib, editors. Handbook of Brain Theory and Neural Networks. 2nd ed. Cambridge, MA: MIT Press; 2002. p. 418-421.
3. Mitri, Floreano, Keller. The evolution of information suppression in communicating robots with conflicting interests. PNAS. 2009; 106:15786-15790.
4. Surfdaddy Orca. Darwin's Robots [document on the Internet]. h+ Magazine; 2009 [cited 2009 November 5]. Available from: http://hplusmagazine.com/articles/ai/darwins-robots
5. CC-BY-NC, Jenn and Tony Bot. Available from http://www.flickr.com/photos/ittybittiesforyou/2275017292/