

Sara Mitri

The Unexplored Territory of Consciousness

June 2003

“The door has been opened to so many new possibilities, questions and fears that can only be overcome by venturing through the door, down the path of the unknown, which is precisely what I intend to do.” —Sara Mitri

Ever since the dawn of philosophy, ever since man started to wonder about life and its meaning, one question has arisen, which has yet to be put to rest. Consciousness is one of the oldest unsolved mysteries of human existence. Answers have been proposed since Socratic times till today; philosophers simply cannot reach a consensus.

Today, the term "consciousness" is still as vague as ever. A real definition does not yet exist. For Marvin Minsky¹, one of the founders of Artificial Intelligence, "to be conscious is to know how you are thinking something". According to his definition, human beings are not really conscious most of the time, as we "have little ability to sense what happens within and outside ourselves." [1] In contrast, Professor Steven Pinker, a cognitive neuroscientist at MIT, describes it as "pure being, subjective experience, raw feelings, first-person present tense, what it is like to see red or feel pain or taste salt." [1] This latter description would correspond to what most people would call "consciousness" and is therefore the meaning I will stick to for the rest of this essay.

Understanding consciousness involves a great number of phenomena that require explanation. As David Chalmers², a Professor of Philosophy at the University of Arizona, sees it, the problems to be tackled can be divided into "easy problems" and "hard problems". The easy problems are ones that can be explained to a great extent using our current knowledge in the field of cognitive sciences, such as the ability to integrate information, to react to external stimuli, to focus the attention or to tell the difference between wakefulness and sleep. The harder part of the problem is how to gain an understanding of the deeper ideas of experience, which include perceptions and emotions. [2]

Up to this point, it may seem that this essay is taking yet another shot at explaining consciousness from a philosophical perspective. However, what interests me as a student of Computer Science is an attempt at solving the ancient problem through the relatively new field of what has been referred to as "Artificial Life", also commonly known as "A-Life".

Coined in 1987 in Los Alamos, New Mexico, "Artificial Life is devoted to the creation and study of lifelike organisms and systems built by humans. The stuff of this life is non organic matter, and its essence is information: computers are the kilns from which these new organisms emerge. Just as medical scientists have managed to tinker with life's mechanisms in vitro, the biologists and computer scientists of A-Life hope to create life in silico." [3] Researchers in artificial life are concerned with modeling, copying or adapting systems from biology. The hope—and indeed, the results that many experiments have shown—is that through imitating biological processes, the complexity of life as we know it in the "real" world will emerge into its silicon model.

One of the more memorable pieces of research in this area was performed by Karl Sims. His work involved creating an artificial universe with its own rules of gravity, friction, light and water, in which he allowed virtual creatures to evolve. The experiment started with individuals possessing very little brain or body, but soon developed, through evolution, creatures with the ability to swim like turtles, or to spiral like snakes through the water. Many of the evolved creatures showed movement schemes alarmingly close to those familiar to us in nature.

Emergent behavior constitutes the basic philosophy behind A-Life. As Sims' research shows, all that was needed was fertile ground for his creatures, an environment to "live" and expand in and a basic structure of a body and a brain. Their life-like behavior is evolved through hundreds of generations of creatures, competing for survival and reproducing offspring accordingly. As Sims sees it, "his virtual creatures had neural networks so complicated that just drawing one of them would require a room-sized sheet of paper, and understanding them would be next to impossible." [1]

Another of the most successful A-Life attempts is the game CREATURES, designed by Steve Grand. CREATURES is a computer program in which one can breed communities of little, furry, virtual animals as if they were pets on the computer. Creatures can be taught speech and manners. The user is responsible for caring for his or her creatures, playing with them, making sure they eat healthily and taking care of them when they get sick. When the creatures go through the stages of puberty, they can be encouraged to mate and to produce offspring. Eventually, the creatures get old and die. The user can then continue raising new generations of creatures. Quite an impressive number of CREATURES enthusiasts now swap their creatures over the Internet, creating a huge community of creatures with different personalities and appearances. According to Grand, there are probably around a million of these creatures in existence at one time, and the gene pool is constantly increasing. Grand admits that he doesn't believe his creatures are actually conscious; instead he refers to them as "quasiconscious". "As systems, they are the recipients of information about their environment (they are 'aware' of the state of the world). They also have information about their own internal state and their actions, so they are 'selfaware' too. But if you knock on their door, I think you'll find there is nobody home." [4] Grand does believe in the possibility of consciousness emerging from artificial life- forms such as his creatures, though. We are just not quite there yet.

In his book, *Creation: Life and How to Make It*, he describes how the emergent properties of Artificial Life can lead us to creating consciousness:

"If we simply try to write a computer program that behaves as if it is conscious, then I think it is wrong to say that the result really is conscious. I do not even think this approach can succeed, in fact, because consciousness is such a complex phenomenon that building its outward properties in a piecemeal way

would be a hopelessly protracted quest, whereas if we can get consciousness to emerge by itself, we will get all of its properties for free." [4]

From a philosophical point of view, I personally believe that what we call "brain", "soul" or "consciousness" are phenomena that we have acquired through evolution. The sheer complexity of the mechanics of our brain has created something that is more than the sum of its parts. I do not believe that there exists a part of the brain that is responsible for consciousness and that can be understood in scientific terms; consciousness is completely detached from the material world, or what philosophers like to refer to as "matter".

That is the reason why creating consciousness should theoretically be an attainable goal. We do not need to build a conscious brain, but to choose a point in the evolution of the brain that we can model sufficiently, as well as the right conditions for it to evolve. The difficulty therefore lies in finding an elegant minimal structure from which to obtain lifelike or hopefully, even conscious behaviors.

Another baffling question that poses itself at this point is how to tell whether a creature is conscious or not. Can we really determine whether a worm, for example, is conscious? Or perhaps, a dog, or even another human being? Many philosophers have put a lot of thought into this matter. René Descartes' famous "I think, therefore I am" [5] represents precisely that idea. The only thing that is certain is one's own existence, one's own consciousness.

Nevertheless, none of the leading scientists in the field seem to doubt the fact that "it is possible"—that one day, sooner or later, whether in a decade, a century or a millennium, the theory seems to confirm the feasibility of creating life, creating consciousness.

Does this then mark the end of our carbon-based life? Does this mean that we are on the verge of creating our own successors in the endless chain of evolution? On the other hand, some see it as the beginning of human immortality. If these new beings will be able to accommodate human consciousness, would it not be possible to download one's own consciousness into another such life-form and live indefinitely in a virtual world? [3] Many of the great thinkers of our age would consider this crossing the line of morality, venturing into dangerous territories that might have uncontrollable consequences. What if the whole situation gets out of hand? Even if it is possible to create life, can we really pull the plug whenever we want to? History has proven time and time again that the advance of science simply cannot be restrained. Human curiosity about this question has not rested since it was posed. Now that we can see a way to answer the question, I doubt that anything will be able to stop it from happening. The door has been opened to so many new possibilities, questions and fears that can only be overcome by venturing through the door, down the path of the unknown, which is precisely what I intend to do.

Notes:

1 Marvin Minsky has made many contributions to AI, cognitive psychology, mathematics, computational linguistics, robotics, and optics. His conception of human intellectual structure and function is presented in *The Society of Mind* (1987), which is also the title of the course he teaches at MIT.

2 Dr. Chalmers is the author of *The Conscious Mind: In Search of a Fundamental Theory* (Oxford University Press, 1996). His article "Facing Up to the Problem of Consciousness" forms the basis of *Explaining Consciousness: The Hard Problem* (MIT Press, 1997, Shear, J. ed.)

References

[1] Bentley, Peter J. *Digital Biology*. London: Headline Book Publishing, 2001.

[2] Chalmers, David. "Facing Up to the Problem of Consciousness." *KurzweilAI.net* (17 August 2002).
Permanent link: <http://www.kurzweilai.net/meme/frame.html?main=/articles/art0512.html>

[3] Levy, Steven. *Artificial Life: The Quest for a New Creation*. New York: Pantheon Books, 1992.

[4] Grand, Steve. *Creation: Life and How to Make It*. London: Weidenfeld and Nicolson, 2000.

[5] Levenson, Carl, and Jonathan Westphal, eds. *Reality*. Hackett Readings in Philosophy. Indianapolis: Hackett Publishing Company, Inc., 1994.