USING ARTIFICIAL LIFE TO ASSESS THE TYPICALITY OF TERRESTRIAL LIFE

M. L. Lupisella

NASA Goddard Space Flight Center, Greenbelt Rd, Greenbelt, MD 20771
University of Maryland, College Park, Department of Biology

ABSTRACT

The extent to which extraterrestrial life questions can be confidently addressed rests in large measure on the extent to which terrestrial life is representative of life in general since we will have to draw from terrestrial life knowledge. This paper outlines a long-term research program that could inform the extent to which terrestrial life is representative of life more generally, which might then help inform our level of confidence in applying terrestrial life knowledge to extraterrestrial life issues. The approach involves appealing to the relatively new field of Artificial Life to: (1) use minimal characterizations of life in (2) a large number of open-ended Artificial Life computer experiments to generate "life possibility spaces" (3) the results of which can be examined for their plausibility within the context of relevant constraining knowledge, so that (4) the remaining results can be examined for variability relative to terrestrial life, where low variability might suggest that terrestrial life is typical of life in general, and high variability could be interpreted to suggest that terrestrial life might be atypical.

INTRODUCTION

This paper will suggest an approach that in the absence of extraterrestrial life could inform the extent to which terrestrial life is representative of life more generally. This could then inform the level of confidence we might have in applying our knowledge of terrestrial biology and ecology to extraterrestrial life issues such as search and detection strategies as well as the interaction of terrestrial ecosystems with other possible planetary ecosystems. The approach involves appealing to the relatively new field of Artificial Life (A-Life) to: (1) use what might be the most minimal set of life-defining characteristics as the basis for (2) a large number of open-ended Artificial Life computer experiments to generate "life possibility spaces", (3) which can be examined for their plausibility within the context of relevant constraining knowledge, so that (4) the remaining possibility space(s) can be examined for variability relative to terrestrial life, where low variability might suggest that terrestrial life is not an anomaly, but is instead sufficiently representative of life in general. High variability in the possibility space(s) could be interpreted to suggest otherwise.

Definitions of Artificial Life

Chris Langton, the first to use the term, "artificial life", suggests biology has traditionally started from the top and "worked *analytically* down from there through the hierarchy of biological organization" (Langton, 1996), where 'analytically' implies the separation of a whole into sub-elements which can be studied individually. This approach seems to have provided a fairly broad picture of the *mechanics* of life on Earth, but Langton suggests that the *dynamics* of life have largely gone unexplored because dynamics is concerned with the interactions *between* parts, which disappear when isolating parts for investigation. Systems with such strong interaction dependency are thought to be non-linear and to require the synthesis of systems to form a coherent whole in order to understand the suite of interactions and how they give rise to overall system behavior. This, according to Langton, is what A-Life attempts to do. This is often accomplished through simulations based on genetic algorithms which are computer

algorithms/simulations in which populations of candidate solutions to a problem are stochastically selected, recombined, mutated, and then either eliminated or retained, based on their relative fitness, where fitness is defined by a fitness function against which the effectiveness (i.e. fitness) of any given solution is measured.

Langton's general conception of A-Life also includes "wetware" which involves wet-bench lab techniques and experiments that use real natural life components in a laboratory environment to direct an artificial evolutionary process toward the production of other real natural life elements such as RNA (Taylor and Jefferson, 1991). However, in trying to assess the degree of universality of terrestrial life, wetware may not be the best approach because of its dependency on what are arguably quite specific biochemical configurations and processes.

Margaret Boden, drawing from much of Langton's thinking, emphasizes A-Life as a field which "uses informational concepts and computer modeling to study life in general, and terrestrial life in particular" (Boden, 1996). Thalmann and Thalmann (1994) write: "A-Life refers to all the techniques that try to recreate living organisms and creatures by computer."

Promising Artificial Life Characteristics

This section will outline a number of key features of A-Life, such as dynamic fitness, the emergence of ecological dynamics, selection for self-reproduction, and open-ended evolutionary outcomes, that make it a promising theoretical approach for understanding biology more generally.

Dynamic Fitness

Formalizing the selection criteria via a program, which itself is allowed to evolve by the co-evolutionary process noted below, can get us close to natural selection by eliminating the a priori, externally imposed selective criteria usually involved in simulation programs. Danny Hillis used selective processes, and more importantly, from work based on the co-evolution of hosts and parasites, allowed for evolving evaluation functions (fitness functions) to efficiently find optimal sorting circuit designs. Instead of having the sorting networks tested against a fixed set of fitness evaluations (i.e. in this case, sorting problems) the sorting problems were allowed to change over time in response to the sorting networks. This prevented the sorting networks from getting stuck on local fitness maxima.

Essentially, these coupled populations, co-evolving via Darwinian selection, can bootstrap each other up the evolutionary ladder far more efficiently than they can alone. Indeed, Hillis' (1991) evolving "computational selective agent", or fitness tests in the form of evolving sorting problems, managed to generate a better design than other well-designed sorting networks. This example demonstrates that fitness can be simulated as a relative, changing quantity, which depends on the details of the system's evolving selective criteria at any given time—just as we see in nature. It also shows the power of using a co-evolutionary approach, and more generally, suggests that intentional design efforts by humans to create optimal systems, including alternative biochemistry, could fall short compared to the efficiency and creativity of automated computational open-ended evolutionary selective processes.

Ecological Dynamics

Computational ecologies get A-Life closer to nature's complex evolutionary dynamics by incorporating many species of organisms co-evolving to form ecological webs. A specific example involves A-Life work based on the game theoretic model, the Iterated Prisoner's Dilemma (Lindgren and Nordhal, 1991). Over the long run, an individual's score is maximized by cooperating, and this cooperative pattern has been shown to emerge via ordinary Darwinian mechanisms such as assuming that individuals want to maximize their immediate pay-off (Hamilton 1981; Axelrod, 1984). The evolutionary and ecological relevance is apparent when we note that strategies were allowed to evolve via an open-ended process by basing the decision on whether or not to cooperate on varying history lengths of previous interactions. The emergence of cooperation supports the suggestion that A-Life experiments are approximating evolution since cooperation has evolved via natural selection on earth.

Selection for Self-Reproduction

Tom Ray took the key step in removing all externally imposed selection criteria when he created his *Tierra* simulation system (Ray, 1996). In this approach, self-reproducing programs compete for computer processing time and memory space where the selective criteria of the programs is the success of self-reproduction. The programs copy themselves, and those which do it best survive and flourish. The fitness function is contained within the basic function of the replication of the organism itself—as in natural selection.

With noisy replication, i.e. mutations, offspring that behave differently can be generated randomly, and these variants, combined with selection for reproduction, have resulted in quite complex and directly biologically relevant "organisms" and phenomena. For example, a parasite program evolved which exploits the ancestor program (the only program written by Ray) by using the ancestor replication loop to replicate itself. This allowed parasites to copy faster since they do not have to replicate the replication loop. But as we see in nature with viruses and cells, the parasites cannot take over the population to the point of driving the ancestor host to extinction, so a coexistence balance results. Another mutant organism resulted which is immune to the parasite because the immune organism makes it impossible for the parasite to use its self-replicating code and is also able to replicate twice as fast, driving the parasite to extinction.

As a tropical biologist, Ray has recognized other biological phenomena resulting in *Tierra* such as punctuated equilibria, competitive exclusion, symbiotic relationships and cheaters. Ray also reports the evolution of novel self-examination where organisms without an ending template evolved and were still able to calculate their size by using a mid-point of their genome, subtracting it from their beginning template and multiplying it by two. This indicates the evolutionary power of his approach.

Ray notes how much of the evolution of his system is driven not just by physical selective constraints of computer processing time and memory (the analog of the non-biological physical environment) but also by interactions with and adaptations to the biotic environment (e.g. other organisms) which is considered to be the primary force for diversification of organisms. Ray takes this as an encouraging sign that his evolutionary system is behaving consistently with nature.

Life As It Could Be

Ray sees open-ended evolution as a key element of life, and being able to represent organisms as self-contained programs, without predefined rules (or a minimal number of rules) allows behavior to emerge that might not otherwise have been known or predicted. It is this possibility of exploring life as it could be that holds such promise for the suggestions of this paper, and the speed and flexibility of computer-based experiments allow for creative and efficient exploration many possibilities.

Clearly, exploring possibilities for life could also inform our understanding of terrestrial biology, as well. Steven Levy (1992), in his popular book on A-Life, writes, "By simulating a kind of life different from that which we are familiar, A-Life scientists seek to explore paths that no form of life in the universe has yet taken, to better understand the concepts and limits of life itself".

Noting the potential for A-Life to strike a healthy intellectual balance, Daniel Dennett (1991) writes: "Artificial Life…can be conceived as a sort of philosophy—the creation and testing of elaborate thought experiments, kept honest by requirements that could never be imposed on the naked mind of a human thinker acting alone. In short, A-Life research is the creation of prosthetically controlled thought experiments of indefinite complexity."

This possibility for studying for life as it could be makes A-Life a powerful tool, and is the critical enabling feature of A-life for the suggestions of this paper.

AN ARTIFICIAL LIFE APPROACH

Characterizing Life

Defining life has been a notoriously difficult and elusive task. Defying definition might at least suggest the possibility that in an important sense, life may be indefinable (Matthews, 1996). Whether or not life is definable could turn out to be an important issue, but a more practical approach is to consider using several or many working characterizations of life without dwelling on whether or not those characterizations are perfect. In fact, as will be discussed later, a perfect definition of life (or whether such a definition is even possible) is not important for the approach being suggested in this paper. What this approach calls for is the abstraction of sufficiently general, fundamental, and *minimal* characterizations of terrestrial life in the hope that computer experiments based on those characterizations produce life possibility spaces with universally relevant *degrees of variation*.

Open-ended Evolution and Self-replication

Tom Ray has suggested only two themes for defining life: self-replication and open-ended evolution. Open-ended evolution essentially implies unpredictability which is an essential characteristic on John Maynard-Smith's

list of life defining characteristics (Bedau, 1996). For Ray, open-ended evolution of synthetic life should "evolve structures or processes that were not designed in or preconceived by the creator" (Ray, 1996, p. 112). It's noteworthy that Ray's use of open-ended evolution does not necessarily require natural selection, or any particular mechanism, as the vehicle for that evolution. Even if we were to choose one of the two minimal characterizations suggested by Ray, we're still left with important challenges. Why should evolution of any kind be a *necessary* characteristic of life? Similarly, why should "self-replication" be a *necessary* characteristic?

Complexity and Organization

Life is often characterized as being highly complex and organized. But how do we reconcile this when thinking about the origin of life? Are we to think that the initial sparks of life, or the first life-forms, were sufficiently complex? We may not have to because Reggia et al (1993) discovered that self-directed replication could be possible for very simple systems. But if indeed the initial steps in forming life did need to be sufficiently complex systems, perhaps for example, as forms of auto-catalytic networks (Kauffman, 1995), then this suggests the possibility that hitting a certain threshold of organized complexity may make the difference between successfully opposing the Second Law of Thermodynamics or not—on local space and time scales.

Opposing the Second Law

Stuart Kauffman's more recent work, *Investigations*, suggests that an "autonomous agent" is "able to reproduce itself and carry out one or more thermodynamic work cycles" (Kauffman, 2000). This is a minimal definition of life that appeals to the phenomenon of opposing the Second Law of Thermodynamics on local space and time scales.

Chris Adami suggests "Life is a property of an ensemble of units that share information coded in a physical substrate and which, in the presence of noise, manages to keep its entropy significantly lower than the maximal entropy of the ensemble, on timescales exceeding the "natural" timescale of decay of the (information-bearing) substrate by many orders of magnitude" (Adami, 1998). The key notion in this definition is keeping entropy significantly lower than the maximal entropy of the ensemble *on timescales exceeding the "natural" timescale*. It is also significant that this definition does not include self-replication as a necessary characteristic. The definition could be made even more general by removing reference to information, units, and substrate:

> Life is a property of an ensemble of *interactions* which, in the presence of noise, manages to keep its entropy significantly lower than the maximal entropy of the ensemble, on timescales exceeding the "natural" timescale of decay of the system by many orders of magnitude.

This is another example of a very general and minimal characterization of life. Exactly how low the entropy needs to be and how many orders of magnitude longer in time the structure and/or function is maintained are important details that will not be addressed here. The primary point is that this characterization serves as an example of the kind of thinking that could be helpful for the purposes of the approach suggested in this paper.

Using Artificial Life Computer Experiments to Generate Life Possibility Spaces

Artificial life simulations vary widely, which is part of the appeal of the field. Given the characterization of life suggested above, data such as natural relaxation times for ensembles (e.g. collection of molecular structures) would have to be a hard-coded in a computer experiment environment, as would basic chemistry.

Some measure of structure and/or function maintenance would also be required, perhaps as a kind of dynamic fitness function. The preferred aim would be the largest number of experiments possible based on a minimal characterization of life. Preferably, the experiments should be run for the computational equivalent of natural evolutionary times in order to allow for sufficient time to generate possibilities.

Knowing when enough experiments have been performed would be difficult, but we might approximate or derive such a number by considering the number of significant evolutionary events, or key evolutionary junctures, that may have occurred on earth. For example, the number of key evolutionary divergences in the earth's history might be used, either directly or more indirectly, to obtain the number of new experiments required.

If this method seems problematic, we might consider that experiments could be run until sufficiently diverse possibilities were generated. If, after running a reasonable number of experiments, much variation was generated, this might signal that enough experiments had been run to satisfy this approach and hence suggest that terrestrial

4

life may not be indicative of life in general. If, however, this approach continually required an usual number of experiments before something sufficiently new was seen, this could signal that the approach will not likely yield a sufficiently diverse possibility space, perhaps either because there were problems with the experiments (e.g. not sufficiently breaking out of terrestrial biology boundary conditions), or perhaps more importantly, because terrestrial life is indeed indicative of life in general.

The number of experiments may also be constrained by time and computational power. Just how constrained the number of experiments would be could turn out to be an important practical limitation of this approach. But again, if the number of runs is statistically sufficient, it may not be that important, and with today's computational power, many comprehensive experiments should be possible.

Applying Constraining Knowledge to the Possibility Space

This approached is based on the requirement that the computer experiment design and results be properly constrained, so that anomalous results do not mislead assessments of the potential variation (Sullins, 2001). A-Life may be able to produce alternative life possibilities, but so can the human imagination—as the abundance of science fiction wonderfully demonstrates. We also know that while imagination can lead us in fruitful directions, unchecked imagination can also lead us astray. Part of what this approach is striving for is the production of possibilities consistent with physical laws and whatever "laws" of biology are thought to be relevant.

In theory, if the rules and initial conditions are set up properly, then all the possibilities generated should be viewed as plausible alternatives. However, mistakes could be made, sufficient knowledge might not be applied to the rules, or a condition could arise as part of the experiment that has no plausible physical basis. Some results, therefore, might have to be dismissed, but perhaps with unusual caution given that this approach is trying to create alternative biologies.

Significance of Remaining Variance and Practical Relevance

The approach outlined in this paper suggests that if a high number of significantly varied alternative life possibilities emerge from properly constrained open-ended evolutionary computer simulations, then we might consider that terrestrial life is not necessarily sufficiently representative of life in general, because such a varied range of alternative life possibilities could be physically realized the universe. Research based on terrestrial biology as applied to extraterrestrial life issues should then be understood in this context, namely that terrestrial biology may not be typical, perhaps giving rise to caution regarding the confidence we have in applying that terrestrial biology knowledge to extraterrestrial life issues. This could affect planetary protection issues such as forward contamination associated with a human mission to Mars (Lupisella, 1999, 2000). It may also suggest the need for more novel, more efficient, and more varied search and detection strategies such as a cooperative robotics approach which might use many varied single function detection agents that can cover a large area, and that can cooperate to apply the fullest range of functionality needed at any given location of interest that is detected by any other robot(s) (Lupisella, 2002). High variability may also suggest the need for rigorous longer-term contingency planning given that detecting and positively identifying an alternative life form could prove much more elusive and require much caution.

The opposite might also be implied in that if novel alternative biologies do not readily emerge, or occur in small numbers with limited variation, then perhaps we might have increased confidence in applying terrestrial biology knowledge to extraterrestrial biology issues. Central to these assessments is the notion of similarity (or dissimilarity) which would have to better defined, preferably quantitatively, in order to assess the extent to which certain outcomes are similar or dissimilar to terrestrial life and to each other. Komosinski et al. (2001) have begun to try to quantify dissimilarity between artificial creatures, so it would appear possible to assess similarity between artificial and real life-forms, qualitatively as well as quantitatively. Such an assessment might be measured at many levels such as molecular, biochemical, and structural deviations from DNA, a more primitive RNA molecule such as 16sRNA, or an RNA-like molecule thought to be a plausible theoretical precursor. If enough experiments were run, degrees of reproducibility may help quantify the probability of outlier results, and patterns among different simulation results might emerge, perhaps then pointing to possible methods for quantifying similarity, making the task of measuring similarity less difficult than it might first appears.

The possibility spaces would have to be analyzed to see if qualitative or quantitative categorizing could be done to play a role in assessing the variability of the simulation results. For example, there may be 100 different possibilities that arise from 100 different simulations, but some, or many, of these may be sufficiently similar so as

to be grouped together into one category or subcategories. A distribution of possibilities would help assess the overall variability of the possibility spaces if such categorization seemed reasonable. For example, a terrestrial life category would be an obvious choice, leaving the variability of the remaining possibility space to analyze. These assessments will not likely be easy to make, and would of course not be determinative, but hopefully informative at least at a macro level.

ANALYSIS

Characterizing Life

Perhaps the most difficult problem with this approach is appropriately characterizing life. Some of these difficulties have been touched on, but there are potentially deeper concerns.

The first concern is a kind of circularity problem based on the fact that a characterization of life must be derived from characteristics that are indicative of terrestrial biology. Any such basic conception of life would be based on our terrestrial knowledge, and hence, arguably, our terrestrial biases, since we have no other physical data. Yet we are trying to transcend terrestrial biases. This "one data point" problem—the one data point being all terrestrial life which is part of the same phylogenetic tree—is a serious epistemological challenge. However, as indicated above, it does seem possible to establish quite fundamental characterizations of life and use them as the foundations on which to run A-Life computer experiments. If the characterizations are sufficiently minimal, terrestrial biases will be somewhat mitigated by producing broad possibility spaces. It is imperfect for sure, but the power of A-Life simulations to generate possibilities may be sufficient to make the assessments this paper suggests.

A deeper challenge is that life may be intrinsically, or at least epistemologically, indefinable. Defining life could be a deep metaphysical and/or epistemological problem in the sense that (1) there may be no clear boundary between life and non-life, (2) we don't know enough to make a clear distinction, or (3) that while there may be a universal characterization, we have our own insurmountable epistemological limitations that would prevent us from obtaining a complete and indisputable characterization of life.

If life is indefinable, then claims of universality are suspect. But, since the approach being suggested in this paper is about asking whether *potentially realizable* life possibilities are *sufficiently varied*, the universality claim has primarily to do with whether that *degree of variance* in the possibility space is universally relevant—that is, that it is possible for the universe to generate that level of variety—not whether the specific results of each experiment outcome are universal. This is one of the main reasons for trying to establish the most minimal characterizations of life and appropriately applying constraining knowledge on the possibility spaces—to avoid unduly restricting results while still maintaining the relevance of the possibility spaces and any associated variance.

One specific practical way to deal with this problem in terms of the approach suggested in this paper is to use many minimal characterizations of life in many different series of experiments, each series or life possibility space being based on its own minimal life characterization, to generate many life possibility spaces, each of which could be analyzed in isolation and relative to each other.

A-Life Computer Experiments vs. Reality

A justified and pervasive concern about models and computer simulation/experiment results is their relation to physical reality. Because the A-Life simulations suggested in this paper are of an open-ended nature, we might think we should be especially concerned about the extent to which the simulation results, i.e. the life possibility spaces, are consistent with physical reality. However, as mentioned in the previous section, it is not the claim of this paper that the simulation results will represent actual or future natural biological phenomena. The suggestion here is only that if the computer experiments produce a sufficiently varied possibility space, that universe may also be capable of generating such variation, and that as a result, we should take seriously the possibility that terrestrial life may be only one of many kinds of life, perhaps calling into question the confidence we might have in applying terrestrial biological knowledge to extraterrestrial life issues.

This suggestion will be supported based primarily on the extent to which this approach can produce *optimal* life possibility spaces, where 'optimal' means that the possibilities are physically plausible and that the emergent variation is indicative of potentially realizable variation in the universe. That is, the experiments may produce results that may or may not ever occur in the universe, but that *could* occur (and whether they do in reality could depend entirely on random factors). Similarly there might be other possibilities and actual natural realizations not represented in the possibility spaces. It is also possible that many or all of the results might never be naturally

6

realized in the universe because of many factors, including random factors—a cornerstone of evolution as we understand it.

Is this Approach Feasible?

There are many "ifs": (1) If appropriately minimal characterizations of life can be established, and (2) if chemistry can be properly programmed in an open-ended evolutionary simulation environment and (3) if experiments based on those characterizations can be run to produce a optimal life possibility spaces, and (4) if we can properly analyze that possibility space for variance, and (5) if we have some idea of the threshold variance that would signal calling into question the confidence we might have that terrestrial life is sufficiently representative of life in general, then we could make at least a first order assessment as to the extent to which terrestrial biological research can be confidently applied to extraterrestrial life issues. Even if all of the above are possible, there is no logical necessity that the artificially produced variance will necessarily be indicative of the kind of variance the universe is capable of, and this will be a difficult open question for broader community to bear in mind. However, for the purposes of a first order assessment in the absence of a physical example of extraterrestrial life, this approach may be worth considering.

The characteristics of Artificial Life covered previously in this paper suggest that Artificial Life has the promise to produce artificially generated physically relevant alternative biologies[1], where presently we only have one data point, namely terrestrial biology—a situation we may be in for a long time. Indeed, part of the motivation for the approach in this paper is to ensure we take proper precautions in maintaining the scientific integrity of that "second data point". As is often the case in science, a broad theoretical framework can facilitate a better understanding of terrestrial life, and perhaps more importantly, if other kinds of life are possible, a better understanding of life in general.

This paper has tried to show that while the Artificial Life methodology presented here presents significant challenges and would likely be a long-term research program, there appear to be no fundamental theoretical barriers preventing the potential of this approach from being realized.[2] Despite the theoretical and speculative nature of this kind of research program, the continued absence of biology's "second data point" suggests that such a research program should be considered to help inform issues regarding possible extraterrestrial life.

ACKNOWLEDGEMENTS

REFERENCES

Adami, C., *Introduction to Artificial Life*, Springer-Verlag, Santa Clara, 1998.

Axelrod, R., *The Evolution of Cooperation*, Basic Books, New York, 1984.

Bedau, M., The nature of life, in *The Philosophy of Artificial Lif*e, edited by M. Boden, Oxford University Press, Oxford, pp. 332-357, 1996.

Boden, M., The intellectual context of artificial life, in *The Philosophy of Artificial Life*, edited by M. Boden, Oxford University Press, Oxford, pp. 1-35, 1996.

Dennett, D., Artificial life as philosophy, in *Artificial Life: An Overview*, ed. by C. Langton, The MIT Press, Cambridge, pp. 291-292, 1991.

Hamilton, W., The genetic evolution of social behavior, *Journal of Theoretical Biology*, **7**, 1-52, 1964.

Hillis, D., Co-evolving parasites improve simulated evolution as an optimization procedure, in *Artificial Life II*, edited by C. Langton, C. Taylor, J. Farmer, S. Rasmussen, (Santa Fe Institute Studies in the Sciences of Complexity, Proceedings, **10**), Addison-Wesley, Redwood City, CA, pp. 313-324, 1991.

Kauffman, S., *At Home in the Universe*, Oxford University Press, New York, 1995.

Kauffman, S., *Investigations*, Oxford University Press, New York, 2000.

---

[1] There are some, such as Tom Ray, who advocate "strong A-Life", which claims A-Life is "real" life—literally alive.

[2] Rodney Brooks (MIT) independently suggested the basic thesis of this paper while we were both in attendance at the American Association for the Advancement of Science "What is Life?" workshop (March 2001, Washington, D.C.).

Komosinski, M., G. Koczyk, M. Kubiak, On estimating similarity of artificial and real organisms, *Theory in Biosciences*, 120, (3-4), 2001.

Langton, C., Artificial Life, in *The Philosophy of Artificial Life*, edited by M. Boden, pp. 39-94, 1996.

Levy. S., *Artificial Life: A Report From Where Computers Meet Biology*, Vintage Books, New York, 1992.

Lindgren, K., and M. Nordhal, Cooperation in artificial ecosystems, in *Artificial Life: An Overview*, edited by C. Langton, The MIT Press, Cambridge, pp. 15-37, 1991.

Lupisella, M., Ensuring the scientific integrity of possible martian life, paper IAA-99-IAA.13.1.08 presented at the International Astronautical Federation Congress, Amsterdam, 1999.

Lupisella, M., Humans and Martians, *Earth Space Review*, **9**, pp. 50-60, 2000.

Lupisella, M., Cooperative Agents and the Search for Extraterrestrial Life, NASA Goddard/JPL Workshop on Radical Agent Concepts proceedings, in press, *Springer-Verlag Lecture Notes in Artificial Intelligence*, 2002.

Matthews, G., Aristotle on Life, in *The Philosophy of Artificial Life*, edited by M. Boden, Oxford University Press, Oxford, pp. 303-313, 1996.

Ray, T., An approach to the synthesis of life, in *The Philosophy of Artificial Life*, edited by M. Boden, Oxford University Press, Oxford, pp. 111-145, 1996.

Reggia, J., S. Armentrout, H. Chou, Y. Peng, Simple systems that exhibit self-directed replication, *Science*, **259**, 1282-1287, 1993.

Sullins, J., Knowing life: possible solutions to the practical epistemological limits in the study of artificial life, *Journal of Experimental & Theoretical Artificial Intelligence*, **13**, pp. 397-408, 2001.

Taylor, C., and D. Jefferson, Artificial life as a tool for biological inquiry, in *Artificial Life: An Overview*, edited by C. Langton, The MIT Press, Cambridge, pp. 1-13, 1991.

Thalmann, N., and D. Thalmann, Creating artificial life in virtual reality, in *Artificial Life and Virtual Reality*, John Wiley and Sons, New York, pp. 1-10, 1994.

Mark Lupisella email address: Mark.L.Lupisella@nasa.gov