

2. The logical Mechanisms of Life

“The designs found in nature are nothing short of brilliant, but the process of design that generates them is utterly lacking in intelligence of its own”. Daniel Dennett [2005]

Life-As-It-Could-Be: but, what is non-life-as-it-could-be? Or how Artificial Life is always theoretical.

“Artificial Life [AL] is the study of man-made systems that exhibit behaviors characteristic of natural living systems. It complements the traditional biological sciences concerned with the *analysis* of living organisms by attempting to *synthesize* life-like behaviors within computers and other artificial media. By extending the empirical foundation upon which biology is based *beyond* the carbon-chain life that has evolved on Earth, Artificial Life can contribute to theoretical biology by locating *life-as-we-know-it* within the larger picture of *life-as-it-could-be*. [Langton, 1989, page 1]

[AL] views life as a property of the *organization* of matter, rather than a property of the matter which is so organized. Whereas biology has largely concerned itself with the material basis of life, Artificial Life is concerned with the formal basis of life. [...] It starts at the bottom, viewing an organism as a large population of *simple* machines, and works upwards *synthetically* from there — constructing large aggregates of simple, rule-governed objects which interact with one another nonlinearly in the support of life-like, global dynamics. The ‘key’ concept in AL is *emergent behavior*.” [Ibid, page 2]

“Artificial Life is concerned with tuning the behaviors of such low-level machines that the behavior that emerges at the global level is essentially the same as some behavior exhibited by a natural living system. [...] Artificial Life is concerned with generating lifelike behavior.” [Ibid, pp 4 and 5]

The previous quotes indicate the goals of Artificial Life according to Chris Langton: the search for complex, artificial, systems which instantiate some kind of lifelike *organization*. The field is interested in both synthesizing an actual artificial living organization, as well as simulating lifelike behavior. The first goal is more ambitious and related to the first definition of life introduced in chapter 1, while the second goal is related to the second definition. The methodology to reach either of these goals is also in line with the notion of emergence mentioned in chapter 1: from the non-linear interaction of simple, mechanistic, components, we wish to observe the emergence of complicated, life-like, unpredictable, behavior.

Natural living organisms are likewise composed of non-living components. As pointed out in chapter 1, the origin problem in biology is precisely the emergence of life from non-living components. The material components of course follow physical law. However, as discussed in chapter 1, a mechanistic explanation of the overall living system is incomplete without addressing biological function and historical context which emerge under natural selection [Gould, 1984; Pattee & Raczaszek-Leonardi, 2012; Rocha, 1998; Laubichler & Renn, 2015; Laubichler et al, 2015]. Similarly, in Langton’s conception, Artificial Life is composed of formal components obeying a particular set of axioms (mechanisms), and from their interaction, global behavior emerges which is not completely explained by the local formal rules. Clearly, the formal rules play the role of an artificial matter and its laws, while the emerging global behavior, if recognized as life-like, plays the role of an artificial biology.

“Of course, the principle assumption made in Artificial Life is that the ‘logical form’ of an organism can be separated from its material basis of construction, and that ‘aliveness’ will be found to be a property of the former, not of the latter.” [Langton, 1989, page 11]

The idea – a systems thinking idea as discussed below – is that if we are able to find the basic design principles of the living organization, then the material substrate used to realize life is irrelevant. By investigating these basic principles we start studying not only biological, carbon-based, life – life-as-we-know-it – but really the general (even universal) rules of life, or life-as-it-could-be. Moreover, from a better understanding of the design principles of life, we can use them to solve engineering problems similar to those that living organisms face [Segel and Cohen, 2001; DeCastro and Von Zuben, 2005]. Several problems have been raised regarding this separation of matter from organization (or form); that is, the search for a universality without matter [Cariani, 1992; Moreno et al, 1994], which will not be discussed here to focus on the relationship between the two distinct goals of AL.

The two goals of AL are usually described as *hard* and *soft AL* respectively. The first concerns the synthesis of artificial life from computational or material (e.g. embodied robotics) components. The second is interested in producing life-like behavior and is essentially metaphorical. Studying and characterizing the collective behavior that emerges from the complex interaction of formal components in search of *interesting* behavior leads to a certain circularity. If AL is concerned with finding (synthesizing or simulating) life-like behavior in artificial systems, we are ultimately binding life-as-could-be to the behavior of life-as-we-know-it by virtue of some subjective resemblance. This can hardly be accepted as a *scientific* search for general principles since experiments are not performed on an independent, natural reality, but rather on artificial components built according to the theory (or design principles) one is testing to begin with: simulations validated by (metaphorical) theory rather than direct experiments.

“They say, ‘Look, isn’t this reminiscent of a biological or a physical phenomenon!’ They jump in right away as if it’s a decent model for the phenomenon, and usually of course it’s just got some accidental features that make it look like something.” [Jack Cowan as quoted in *Scientific American*, June 1995 issue, “From Complexity to Perplexity”, by J. Horgan, page 104]

“Artificial Life — and the entire field of complexity—seems to be based on a seductive syllogism: There are simple sets of mathematical rules that when followed by a computer give rise to extremely complicated patterns. The world also contains many extremely complicated patterns. Conclusion: Simple rules underlie many extremely complicated phenomena in the world. With the help of powerful computers, scientists can root those rules out.” [J. Horgan, *Scientific American*, June 1995 issue, “From Complexity to Perplexity”, page 107]

“Artificial Life is basically a fact-free science”. [John Maynard Smith as quoted in *Scientific American*, June 1995 issue, “From Complexity to Perplexity”, by J. Horgan, page 107]

The problem is that AL models and theory must be experimentally validated against something, otherwise it becomes a fact-less manipulation of computer rules with subjective resemblances to biology. AL produces formal systems that display emergent complex behaviors, but which behaviors can be classified as “life-as-it-could-be” and which cannot? What is the formal *threshold of complexity* needed? In the natural world we are able to distinguish life from non-life, biology from physics due to the known signatures of bio-chemistry. Furthermore, we can experiment with biology directly, since it is materially available independently of any theory we may have about it. But in the formal realm, we need additional criteria to distinguish logical life from logical non-life, artificial life from artificial physics. Since the criteria are to be tied to life-as-it-could-be, per Langton’s framing, they must ultimately be theoretical—i.e. based on our theories of life, not on life itself. Life-as-it-could-be is never available for direct experimentation; scientific experimentation requires implementation of some theoretical model (simulations), even if we engineer them materially.

“Artificial Life must be compared with a real or an artificial nonliving world. Life in an artificial world requires exploring what we mean by an alternative physical or mathematical reality.” [Pattee, 1995]

In a strict sense, thus, AL cannot be an experimental science and must remain theoretical by definition. Its methodology requires existing theories of life to be compared against. Still, it can certainly contribute to the meta-methodology of Biology by allowing us to test and improve biological theory beyond unavoidable material constraints, such as the incomplete fossil record or fast measurement of cellular activity. From this point of view and to be useful, AL should not settle for vague rules of what constitutes living behavior. Whether we want to synthesize life or merely simulate a particular behavior of living organisms, the field should aim to investigate the clear rules that allow us to distinguish life from non-life in a sound theoretical framework that remains in sync with “life-as-we-know-it” – especially because the reality of living matter has proven to be stranger than our fiction (or theories) of it (see Chapter 5). Only by establishing an artificial physics, from which an artificial biology can emerge, and a *theory*, or set of rules, distinguishing the two, can it explore a generalized theory of life. Naturally, the requirements for hard AL are much stricter, as we are not merely interested in behaviors that can be compared to real biological systems with looser or stricter rules, but the engineering (realization) of an artificial organization that must be agreed to be living according to some clear theory. Soft AL can restrict itself to particular behavioral traits which need only to be simulated to a satisfactory degree.

Simulations, Realizations, Systemhood, Thinghood, and Theories of Life

“Boids are *not* birds; they are not even remotely like birds; they have no cohesive physical structure, but rather exist as information structures — processes — within a computer. But — and this is the critical ‘but’ — at the level of behaviors, *flocking Boids and flocking birds are two instances of the same phenomenon*: flocking. [...] The ‘artificial’ in Artificial Life refers to the component parts, not the emergent processes. If the component parts are implemented correctly, the processes they support are genuine — every bit as genuine as the natural processes they imitate. [...] Artificial Life will therefore be genuine life — it will simply be made of different stuff than the life that has evolved on Earth.” [Langton, 1989, pp. 32-33]

“Simulations and realizations belong to different categories of modeling. Simulations are metaphorical models that symbolically ‘stand for’ something else. Realizations are literal, material models that implement functions. Therefore, accuracy in a simulation need have no relation to quality of function in a realization. Secondly, the criteria for good simulations and realizations of a system depend on our theory of the system. The criteria for good theories depend on more than mimicry, e.g., Turing Tests.” [Pattee, 1989, page 63]

As Pattee points out, the bottom line is that a simulation, no matter how good it is, is not a realization. Nonetheless, it may still be possible to obtain artificial living organisms (realizations) if, from an artificial environment, we are able to generate, in a bottom-up manner, organizations which match some theory of life we wish to test. Howard Pattee [1989] has proposed that if emergent artificial organisms are able to perform measurements, or in other words, categorize their (artificial or natural) environment, then they may be considered realizations. Some claim that computational environments do not allow for this creative form of emergence [see Cariani, 1992; Moreno, et al, 1994]. In any case, whatever environment we may use, computational or material, the very conception of a living phenomenon requires a theory to recognize it—to distinguish life from non-life. For instance, the flocking behavior in Langton’s quote above assumes a theory about flocking such that we can say that flocking boyds and flocking birds are instances of the same phenomenon—and boyds and similar systems are circularly devised so that they meet such a theory.

In contrast, by being circumscribed to organizations that exist in and reproduce and evolve by DNA-RNA-Protein biochemistry, *life-as-we-know-it* can be experimentally studied without a theory of biological phenomena. One can counter-argue that known biochemistry assumes an (inductive) theory that all life is made of DNA-RNA-Protein biochemistry. But since that biochemistry is universal for (recognized) life on Earth, we can work within this universe without considering a larger, encompassing “general theory” of life,

unless and until another life form is discovered. Certainly at the edges of known Biology one can ask if such evolving things as viruses are alive? Since at its most foundational level life is recognized as organizations that reproduce and evolve by DNA-RNA-Protein biochemistry, even the autonomy of such reproduction does not need to be (theoretically) postulated and observed. Thus it is easy to agree that viruses that reproduce in networks of DNA-RNA-Protein biochemistry (that include other organisms) are alive—even if their reproduction is not autonomous.

Certainly theory is useful and indeed foundational for biology, e.g. Natural Selection. But hypothesis validation in the life sciences can focus on experimental mechanism identification without needing to justify that the organisms or behaviors under study constitute biological phenomena. In contrast, in AL, experimenting on (artificial) mechanisms is always tied to a generalized phenomenon such as flocking, which needs to be defined as a guiding theory or general system [Klir, 2001]. Indeed, AL can be seen as a discipline (focusing on life phenomena) of complex systems science, which encompasses the search of those *organizational* properties of the universe which can be abstracted from their specific material substrate: *systemhood* from *thinghood* [Rosen, 1986; Klir, 2001]. The abstraction of general principles of organization is always dependent on a theory which conceives organizations that may never exist in reality in their pure form, but which are nonetheless useful not only for theory but for experimental approximations via models and simulations, e.g. ideal gases and scale-free networks [Barrat, Barthelemy, and Vespignani, 2008; Broido & Clauset, 2019; Holme, 2019]. But the usefulness of general systems, as Pattee's quote stresses, depends on more than vague mimicry. It depends on the accurate prediction of natural phenomena such as climate and pandemics – *as-we-know-them* and can access experimentally, not just *as-they-could-be* theoretically.

The difficulty for systems science, or complexity theory, lies precisely in the choice of the appropriate level of abstraction. If we abstract too much and do not validate against natural phenomena, most things will look alike, leading to a theory of fact-less, reminiscent analogies, exposed by Cowan and Maynard-Smith above. If, on the other hand, we abstract too little, inquiry reduces to increasingly specific niches, accumulating much data and knowledge about (context-specific) components without much *understanding* of, or ability to *control*, the (general) macro-level organization. In the context of life, we do not want to be tied uniquely to knowledge about gene *x* in organism *y*, but we also do not want *life-as-could-be* to be anything at all that makes no predictions about *life-as-we-know-it* or even unknown natural life forms we may find. The challenge lies precisely on finding the right amounts of theoretical *systemhood* and experimental *thinghood*, as well as the interactions between the two, necessary for a good theory of life, real or artificial.

A final unanswered question is whether there are systems from which *systemhood* cannot be completely abstracted from *thinghood*? Life is sometimes proposed as one of those systems [see Rosen, 1986, 1991; Moreno et al, 1994; Pattee, 1995]. However, as we will see in this course (see remaining chapters), we can at least demonstrate that there are many general systems derived from biological systems which lead to very useful bio-inspired computing algorithms, suggesting that systemhood principles can be extracted and indeed help explain life phenomena such as adaptation, evolution, self-organization, control, stigmergy, collective behavior, immunity, and intelligence.

Further Readings and References

- Barrat, A., M. Barthelemy, and A. Vespignani [2008]. *Dynamical processes on complex networks*. Cambridge university press, pp. 756-59.
- Broido, A. D., & Clauset, A. (2019). "Scale-free networks are rare" . *Nature Communications*, **10**(1), 1017.
- Cariani, P. [1992], "Emergence and Artificial Life" In *Artificial Life II*. C. Langton (Ed.). Addison-Wesley. pp. 775-797.
- de Castro, L. N. & Von Zuben, F. J. (eds.) [2005]. *Recent Developments in Biologically Inspired Computing*. Idea Group Publishing.

- Dennet, D.C. [2005]. "Show me the Science". *New York Times*, August 28, 2005.
- Holme, P. (2019). "Rare and everywhere: Perspectives on scale-free networks". *Nature Communications*, **10**(1), 1016.
- Langton, C. [1989], "Artificial Life" In *Artificial Life*. C. Langton (Ed.). Addison-Wesley. pp. 1-47.
- Laubichler, M. D., Stadler, P. F., Prohaska, S. J., & Nowick, K. [2015]. "The relativity of biological function". *Theory in Biosciences*, **134**(3), 143-147.
- Laubichler, M. D., & Renn, J. [2015]. "Extended evolution: A conceptual framework for integrating regulatory networks and niche construction". *Journal of Experimental Zoology Part B: Molecular and Developmental Evolution*, **324**(7), 565-577.
- Klir, G. [2001], *Facets of Systems Science*. Springer. (On Constructivism pp. 12-13)
- Moreno, A., A. Etxeberria, and J. Umerez [1994], "Universality Without Matter?". In *Artificial Life IV*, R. Brooks and P. Maes (Eds). MIT Press. pp 406-410
- Pattee, H. [1989], "Simulations, Realizations, and Theories of Life". In *Artificial Life*. C. Langton (Ed.). Addison-Wesley. pp. 63-77.
- Pattee, H. [1995], "Artificial Life needs a real Epistemology". In *Advances in Artificial Life*. F. Moran, A Moreno, J.J. Merelo, P. Chacon (Eds.). Springer-Verlag.
- Pattee, H. H., & Raczaszek-Leonardi, J. [2012]. *Laws, language and life: Howard Pattee's classic papers on the physics of symbols with contemporary commentary* (Vol. 7). Springer Science & Business Media.
- Rocha, L. M. [1998]. "Selected self-organization and the semiotics of evolutionary systems". In *Evolutionary systems* (pp. 341-358). Springer, Dordrecht.
- Rosen, R. [1986], "Some Comments on Systems and System Theory". In *Int. Journal of General Systems*. Vol. 13, No.1.
- Rosen, R. [1991]. *Life Itself: A Comprehensive Inquiry into the Nature, Origin, and Fabrication of Life*. Columbia University Press.
- Segel, L.A. and I.C. Cohen [2001]. *Design Principles for the Immune System and Other Distributed Autonomous Systems*. Santa Fe Institute Studies in the Sciences of Complexity. Oxford University Press.

For next lecture read:

- Floreano, D. and C. Mattiussi [2008]. *Bio-Inspired Artificial Intelligence: Theories, Methods, and Technologies*. MIT Press. Preface, **Sections 4.1, 4.2.**
- Optional: Nunes de Castro, Leandro [2006]. *Fundamentals of Natural Computing: Basic Concepts, Algorithms, and Applications*. Chapman & Hall. **Chapter 7, sections 7.1, 7.2 and 7.4.**
- Optional: Part I of Flake's [1998], *The Computational Beauty of Life*. MIT Press.