

1. What is Life?

“What was life? No one knew. It was undoubtedly aware of itself, so soon as it was life; but it did not know what it was”. Thomas Mann [1924]

Threshold of Complexity

“Seeking a connecting link, they had condescended to the preposterous assumption of structureless living matter, unorganized organisms, which darted together of themselves in the albumen solution, like crystals in their mother-liquor; yet organic differentiation still remained at once condition and expression of all life. One could point to no form of life that did not owe its existence to procreation by parents”. Thomas Mann [1924].

“Nothing in biology makes sense without evolution”. Theodosius Dobzhansky [1973]

Biologically-inspired computing is an interdisciplinary field that formalizes processes observed in living systems to design computational methods for solving complex problems, or simply to endow artificial systems with more natural traits. But to draw more than superficial inspiration from Biology we need to understand and discuss the concept of *life*. It should be noted that for the most part of the history of humanity, the question of what *life* is was not an important issue. Before the study of mechanics became important, everything was thought to be alive: the stars, the skies, the rivers and mountains, etc. There was no non-life, so the concept was of no importance. It was only when people started to see the World as determined by the laws of mechanics that the question arose. If all matter follows simple physical laws, then what is indeed the difference between life and non-life, between biology and physics? Let us then start with a current dictionary definition:

“life adj.— n.1. the general condition that distinguishes organisms from inorganic objects and dead organisms, being manifested by growth through metabolism, a means of reproduction, and internal regulation in response to the environment. 2. the animate existence or period of animate existence of an individual. 3. a corresponding state, existence, or principle of existence conceived of as belonging to the soul. 4. the general or universal condition of human existence. 5. any specified period of animate existence. 6. the period of existence, activity, or effectiveness of something inanimate, as a machine, lease, or play. 7. animation; liveliness; spirit: The party was full of life. 8. the force that makes or keeps something alive; the vivifying or quickening principle.” [Random House Webster’s Dictionary]

The definitions above fall into three main categories: (1) life as an *organization* distinct from inorganic matter (with an associated list of properties), (2) life as a certain kind of *animated behavior*, and (3) life as a special, incommensurable, quality—*vitalism*. Throughout this course we will see that all principles, and indeed all controversies, associated with the study of life fall into one of these categories or the differences among them. The third category has been discarded as a viable scientific explanation, because for science nothing is in principle incommensurable. The question of whether life is organized according to a special design, intelligent or mysterious, pertains to metaphysics. If the agent of design cannot be observed with physical means, then it is by definition beyond the scope of science as it cannot be measured, and any theories derived from such a concept cannot be tested.

While metaphysical dispositions do not pertain to science, systems thinking has led many scientists to

observed that a naive mechanistic decomposition of life may also fail to explain it [Klir, 2001]. The traditional scientific approach has led the study of living systems into a reductionist search for answers in the nitty-gritty of the biochemistry of living organisms. In this approach, life is nothing more than the complicated physics of a collection of moving bodies. However, the question remains unanswered since there are many ways to obtain some complicated dynamics, but of all of these, which ones can be classified as *alive*? What *kind of complexity* are we looking for? No one disputes that life is some sort of complex material arrangement, but when does an organization reach a necessary *threshold of complexity* after which it is said to be living? Is it a discrete step, or is life a fuzzy concept? Must such an organization have a material implementation (possess a *thinghood*), or can it be completely abstracted (fully explained by its *systemhood*) [Rosen, 1986]? To understand life without meaningless reduction, must we synthesize organizations with the same threshold of complexity (first category above), or is it enough to simulate its animated behavior (second category above) [Pattee, 1989]?

Information Organizes and Breeds Life

“*Life is a dynamic state of matter organized by information*”. Manfred Eigen [1992]

“*Life is a complex system for information storage and processing*”. Minoru Kanehisa [2000]

Traditionally life has been identified with material organizations which observe certain lists of properties, e.g. *metabolism, adaptability, self-maintenance* (autonomy), *self-repair, growth, replication, evolution*, etc. Most living organisms follow these lists, however, there are other material systems which obey only a subset of these rules, e.g. viruses, candle flames, the Earth, certain robots, etc. This often leads to the view that life is at best a fuzzy concept and at worst something we are, subjectively, trained to recognize—life is what we can eat—and is thus not an objective distinction. The modern-day molecular biology view of life, on the other hand, tends to see life as a material organization that if not completely programmed by *genomic information*, is at least very controlled by it [Brenner, 2012, Cobb, 2013]. Thus, when Craig Venter’s team [Gibson et al, 2010] recently produced bacteria with a “*prosthetic genome*” [a term coined by Mark Bedau, see *Nature | Opinion*, 2010] copied from other bacteria but synthesized in the lab, the momentous synthetic biology feat was announced as the creation of the first synthetic or artificial life form.

Indeed, it is remarkable that in Venter’s experiment, a cell with a synthesized prosthetic genome from a similar but distinct organism, was able to reproduce over and over resulting in a cell with a different phenotype from the original, implanted cell—in effect, a cell *re-programmed* by a synthesized genome. Is life then a type of computer that can be reprogrammed [Brenner, 2012]? This also leads us to question how general-purpose can such genomic re-programming be? Will it be restricted to very narrow classes of similar organisms, or will it ever be possible to re-program any prokaryotic or eukaryotic cell ?

The artificial life field, whose members tend to follow the fuzzy list of properties conception of life, does not typically recognize Venter’s bacteria with a prosthetic genome as a bona fide synthesis of artificial life, since it relies on the pre-existence of a working, naturally-obtained cell to implant a prosthetic genome into. Even most molecular biologists will agree that we are nowhere near understanding, let alone synthesizing an artificial cell from scratch [e.g. George Church, see *Nature | Opinion*, 2010]. Nonetheless, Venter’s achievement begs at least the question of what is it about life’s *design principle* that makes it easier to synthesize a working prosthetic genome than a working “prosthetic proteome or metabolome”? It also makes us think about what does “understanding life” mean for biology, biomedical complexity, artificial life, and computation? Why is *genetic information* so important and how does it relate to *information technology*?

Life requires the ability to both categorize and control events in its environment in order to survive. In other words, organisms pursue (or even decide upon) different actions according to *information* they perceive in an environment. Furthermore, living organisms reproduce and develop from genetic *information*. More specifically, genetic information is *transmitted* “vertically” (inherited) in phylogeny and cell reproduction, and *expressed* “horizontally” within a cell in ontogeny for the functioning of living organisms as they interact and react with their environments—we are now sure that genetic information can also be transmitted horizontally between organisms and play an important role in evolution [Goldenfeld & Woese 2007; Riley, 2013]. Indeed, the difference between living and non-living organizations seems to stand on the ability of the former to use *relevant* information for their own functioning. It is this “relevant” which gives life an extra attribute to simple mechanistic interactions. When an *organization* is able to recognize and act on aspects of its environment which are important *to its own* survival, we say that the mechanisms by which the organization recognizes and acts are *functional* in reference to itself (*self-reference*).

We should note that Physics is not concerned with function [Pattee, 1978]. A physical or chemical description of DNA is certainly possible, but will tell us nothing as to the function of a DNA molecule as a gene containing relevant information for a particular organism—in this case, information derives from the sequence of four possible nucleotides not from their minute chemical details [Eigen, 1992]. Only in *reference* to an organism in a particular context does a piece of DNA *function* as a gene (e.g. an enzyme with some biochemical effect in an organism in an environment). The effective separation of the material from the functional and organizational is very much a systems concept, leading many to study alternative material arrangements and general organizations that can sustain life.

Emergence and Explanation

“First, nothing in biology contradicts the laws of physics and chemistry; any adequate biology must be consonant with the ‘basic’ sciences. Second, the principles of physics and chemistry are not sufficient to explain complex biological objects because new properties emerge as a result of organization and interaction. These properties can only be understood by the direct study of the whole, living systems in their normal state. Third, the insufficiency of physics and chemistry to encompass life records no mystical addition, no contradiction to the basic sciences, but only reflects the hierarchy of natural objects and the principle of emergent properties at higher levels of organization”. Stephen Jay Gould [1984].

The issue of function and reference in the living organization is often approached by appealing to the notion of *emergence* or *collective behavior*. Whatever (macro-level) organization exists after the complexity threshold for life is passed, we may say that it is emergent because its attributes cannot be completely explained by the (micro-) physical level. Function, control, and categorization cannot be explained solely by the mechanics and dynamics of the components (individual variables) of life. Understanding and predicting life, requires the study of the *collective behavior* that emerges from complex multivariate, intra- and inter-organism exchanges in eco-social (and even technological) multi-level networks. Notice, however, that emergence does not imply vitalism or dualism. When we say that certain characteristics of life cannot be explained by the physics of biochemical variables alone, we mean that they must be explained by complementary, additional models—namely, *informational*, *historical* and *functional* descriptions. In other words, though biological function, control, and categorization cannot be explained by physics alone, organisms, like anything else, must nonetheless follow physical laws. But information is contextual and historical, and therefore requires more than universal models: it requires contingent, context-specific descriptions. For instance, the *INS* gene encodes instructions to produce *insulin* in humans (and a few other species, with orthologs in hundreds of others). But the same gene is meaningless for the biochemistry of other species, even allowing us to use bacteria as (syntactic) factories of human insulin.

The *origin of life*, is thus a problem of emergence of contextual information from a physical milieu (of universal laws) under specific constraints [Eigen, 1992]. This is the crux of complex systems: the interplay between micro- and macro-level constraints determines their behavior, and multiple, non-decomposable, complementary scales or levels (emergence) are required to understand biological (and biomedical) complexity [Pattee, 1978]. The definition of emergence as an epistemological, explanatory requirement, is related to the notion of *emergence-relative-to-a-model* [Rosen, 1985; Cariani, 1989] or *intensional emergence* [Salthe 1991]. It refers to the impossibility of epistemological reduction of the properties of a system to its components [Clark, 1996]. As an example, we can think of phase transitions such as that of water in its transition from liquid to gas. Water and its properties cannot be rephrased in terms of the properties of hydrogen and oxygen, it needs a qualitatively different model. Another example of complementary models of the same material systems is the wave-particle duality of light.

Physicists *understand* the laws of nature (as best they can), but it takes engineers to *control* nature. The very best physicists are the very best engineers, but those are arguably rare (e.g. Von Neumann). The goal of complex systems science is to understand *organized complexity* (life, society, cognition) in the same way physicists understand nature [Weaver, 1948]. Biology, as a discipline, has not entirely “made up its mind” if it wants to understand life as a physicist or control it as an engineer. Due to its focus on the micro-level of life, its biochemistry, molecular biology follows essentially a (reverse-) engineering, black-box methodology (knockouts, controls, etc.) to understand *mechanism*. This leads to a bit of a schizophrenic agenda: focusing exclusively on micro-level experiments in order to *suggest* macro-level understandings. If the goal is control of biology, say for biomedical advances, then focus should turn to biotechnology engineering—synthetic biology is a good example of this focus shift. If the goal is understanding, then focus should be more on macro-level organized complexity. Ideally, a healthy life sciences program would tie the need to understand complexity with the need to control mechanisms—like physicists and engineers do.

This is where complex systems, artificial life, and bio-inspired computing can contribute to a wider arena of the life sciences; they can be used as laboratories for experimenting with theories of organized complexity, and thus enrich our understanding of life. Artificial life concerns both the simulation and realization of life in some artificial environment, usually the computer. At least regarding the second of its goals, artificial life aims to understand the fundamental micro/macro-level interaction that leads to organized complexity. Bio-inspired computing, as a more pragmatic endeavor, does not need to concern itself with synthesizing actual life, but only with drawing analogies from life (real and artificial). Nonetheless, if the main motivation of bio-inspired computing is that life with its designs has already solved versions of many complex engineering problems we are interested in, then a thorough and accurate understanding of the essential characteristics of life is inescapable. Moreover, by abstracting context-specific principles of life to make them relevant in other settings, provides a useful laboratory to experiment with theoretical biology.

Further Readings and References:

- Brenner, S. [2012]. "Turing centenary: Life's code script." *Nature* **482** (7386): 461-461.
- Cariani, Peter [1989]. *On the Design of Devices with Emergent Semantic Functions*. PhD.Dissertation. SUNY Binghamton.
- Clark, Andy [1996]. "Happy couplings: emergence and explanatory interlock." In: *The Philosophy of Artificial Life*. M. Boden (ed.). Oxford University Press, pp. 262-281.
- Cobb, Matthew. [2013]. "1953: When Genes Became 'Information'." *Cell* **153** (3): 503-506.
- Dobzhansky, T. [1973]. "Nothing in Biology Makes Sense Except in the Light of Evolution". *The American Biology Teacher*, March 1973 (35:125-129)
- Eigen, M. [1992]. *Steps Towards Life*. Oxford University Press.
- D. G. Gibson et al [2010]. "Creation of a Bacterial Cell Controlled by a Chemically Synthesized Genome". *Science*. **329**

(5987): 52-56

- Goldenfeld, Nigel, and Carl Woese [2007]. "Biology's Next Revolution." *Nature* **445** (7126): 369.
- Gould, Stephen Jay [1984]. *Natural History*; Jan84, Vol. 93 Issue 1, p24.
- Klir, G. [2001], *Facets of Systems Science*. Springer.
- Mann, T. [1924]. *The Magic Mountain*. As quoted by Eigen [1990].
- Nature | Opinion [2010] "Life after the synthetic cell". *Nature* **465**: 422–424
- Pattee, Howard H. [1978]. "The complementarity principle in biological and social structures." In: *Journal of Social and Biological Structures* Vol. 1, pp. 191-200.
- Pattee, H. [1989], "Simulations, Realizations, and Theories of Life". In *Artificial Life*. C. Langton (Ed.). Addison-Wesley. pp. 63-77.
- Polanyi, M. [1968]. "Life's irreducible structure". *Science*, **160** (3834), 1308-1312.
- Riley, D. R., K.B. Sieber, K. M. Robinson, J. R. White, A. Ganesan, S. Nourbakhsh, and J. C. Dunning Hotopp [2013]. "Bacteria-Human Somatic Cell Lateral Gene Transfer Is Enriched in Cancer Samples." *PLoS Computational Biology* **9** (6): e1003107.
- Rosen, R. [1986], "Some Comments on Systems and System Theory". In *Int. Journal of General Systems*. Vol. 13, No.1.
- Salthe, Stanley N. [1991], "Varieties of Emergence". *World Futures* Vol. 32, pp.69-83
- Schrödinger, Erwin [1944]. *What is Life?*. Cambridge University Press.
- Weaver, W. [1948]. "Science and Complexity". *American Scientist*, **36**(4): 536-44.

For next lectures read:

- Dennet, D.C. [2005]. "Show me the Science". *New York Times*, August 28, 2005.
- Polt, R. [2012]. "Anything but Human". *New York Times*, August 5, 2012

Optional for ISE483

- Langton, C. [1989], "Artificial Life" In *Artificial Life*. C. Langton (Ed.). Addison-Wesley. pp. 1-47.
- Pattee, H. [1989], "Simulations, Realizations, and Theories of Life". In *Artificial Life*. C. Langton (Ed.). Addison-Wesley. pp. 63-77.

Optional

- Aleksander, I. [2002]. "Understanding Information Bit by Bit". In: *It must be beautiful : great equations of modern science*. G. Farmelo (Ed.), Grant.
- Cobb, Matthew. [2013]. "1953: When Genes Became 'Information'." *Cell* **153** (3): 503-506.
- Gleick, J. [2011]. *The Information: A History, a Theory, a Flood*. Random House. Chapter 8.
- James, R., and Crutchfield, J. (2017). Multivariate Dependence beyond Shannon Information. *Entropy*, **19**(10), 531.
- Nunes de Castro, Leandro [2006]. *Fundamentals of Natural Computing: Basic Concepts, Algorithms, and Applications*. Chapman & Hall. **Chapter 1**, pp. 1-23.
- Prokopenko, Mikhail, Fabio Boschetti, and Alex J. Ryan. "An information theoretic primer on complexity, self organization, and emergence." *Complexity* **15.1** (2009): 11-28.

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 lifelike *organization*. The field is interested in both synthesizing the living organization, and in 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 in line with the notion of emergence also mentioned in chapter 1: from the non-linear interaction of simple, mechanistic, components, the emergence of complicated, life-like, unpredictable, behavior is to be observed.

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 & Racaszek-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 (formal mechanisms), and from their interaction, global behavior emerges which is not completely explained by the (local) component logical rules or axioms. Clearly, the component formalisms 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 indeed 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 boyds in Langton’s quote above assume *a theory about flocking behavior* 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 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 identifying mechanisms experimentally without needing to justify that the organisms or behaviors under study constitute biological phenomena. In contrast, in AL, experiments with (artificial) mechanisms are 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 scientific approach to life, real or artificial, that is both theoretical and experimental, as any scientific approach must be.

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. Etxebarria, 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.

3. Modeling the World and Systems Approach to Life

“When you can measure what you are speaking of and express it in numbers you know that on which you are discoursing. But if you cannot measure it and express it in numbers. your knowledge is of a very meagre and unsatisfactory kind.”. (Lord Kelvin)

The Nature of Information and Information Processes in Nature¹

The word **information** derives from the Latin *informare* (*in* + *formare*), meaning to give form, shape, or character to. Etymologically, it is therefore understood to be the formative principle of something, or to imbue with some specific character or quality. However, for hundreds of years, the word information is used to signify knowledge and aspects of cognition such as meaning, instruction, communication, representation, signs, symbols, etc. This can be clearly appreciated in the *Oxford English Dictionary*, which defines information as “the action of informing; formation or molding of the mind or character, training, instruction, teaching; communication of instructive knowledge”.

When we look at the world and study reality, we see order and structure everywhere. There is nothing that escapes description or explanation, even in the natural sciences where phenomena appear sometimes catastrophic, chaotic and stochastic. A good example of order and information are our roads. Information can be delivered by signs. Drivers know that signs are not distant things, but they are *about* distant things in the road. What signs deliver are not things but a sense or knowledge of things – a message. For information to work that way, there have to be signs. These are special objects whose function is to be *about* other objects. The function of signs is reference rather than presence. Thus a system of signs is crucial for information to exist and be useful in a world, particularly for the world of drivers!

The central structure of information is therefore a relation among signs, objects or things, and agents capable of understanding (or decoding) the signs. An AGENT is *informed* by a SIGN about some THING. There are many names for the three parts of this relation. The AGENT can be thought of as the recipient of information, the listener, reader, interpretant, spectator, investigator, computer, cell, etc. The SIGN has been called the signal, symbol, vehicle, or messenger. And the about-some-THING is the message, the meaning, the content, the news, the intelligence, or the information.

The SIGN-THING-AGENT relation is often understood as a sign-system, and the discipline that studies sign systems is known as *Semiotics*. In addition to the triad of a sign-system, a complete understanding of information requires the definition of the relevant context: an AGENT is informed by a SIGN about some THING in a certain CONTEXT. Indeed, (Peircean) semiotics emphasizes the *pragmatics* of sign-systems, in addition to the more well-known dimensions of *syntax* and *semantics*. Therefore, a complete (semiotic) understanding of information studies these three dimensions of sign-systems:

1. **Semantics**: the content or meaning of the SIGN of a THING for an AGENT; it studies all aspects of the relation between signs and objects for an agent, in other words, the study of meaning.
2. **Syntax**: the characteristics of signs and symbols devoid of meaning; it studies all aspects of the relation among signs such as their rules of operation, production, storage, and manipulation.

¹ This subsection is an excerpt of [Rocha and Schnell, 2005]

3. **Pragmatics**: the context of signs and repercussions of sign-systems in an environment; it studies how context influences the interpretation of signs and how well a signs-system represents some aspect of the environment.

Signs carry information content to be delivered to agents. However, it is also useful to understand that some signs are more easily used as referents than others. In the beginning of the 20th century, Charles Sanders Peirce defined a typology of signs:

1. **Icons** are direct representations of objects. They are similar to the thing they represent. Examples are pictorial road signs, scale models, and of course the icons on your computer. A footprint on the sand is an icon of a foot.
2. **Indices** are indirect representations of objects, but necessarily related. Smoke is an index of fire, the bell is an index of the tolling stroke, and a footprint is an index of a person.
3. **Symbols** are *arbitrary* representations of objects, which require exclusively a social convention to be understood. A road sign with a red circle and a white background denotes something which is illegal because we have agreed on its arbitrary meaning. To emphasize the conventional aspect of the semantics of symbols, consider the example of variations in road signs: in the US yellow diamond signs denote cautionary warnings, whereas in Europe a red triangle over a white background is used for the same purpose. We can see that convention establishes a code, agreed by a group of agents, for understanding (decoding) the information contained in symbols. For instance, smoke is an index of fire, but if we agree on an appropriate code (e.g. Morse code) we can use smoke signals to communicate symbolically.

Clearly, signs may have iconic, symbolic and indexical elements. Our alphabet is completely symbolic, as the sound assigned to each letter is purely conventional. But other writing systems such as Egyptian or Mayan hieroglyphs, and some Chinese characters use a combination of phonetic symbols with icons and indices. Our road signs are also a good example of signs with symbolic (numbers, letters and conventional shapes), iconic (representations of people and animals) and indexical (crossing out bars) elements.

Finally, it is important to note that due to the arbitrary nature of convention, symbols can be manipulated without reference to content (syntactically). This feature of symbols is what enables computers to operate. As an example of symbol manipulation without recourse to content, let us re-arrange the letters of a word, say “deal”: dale, adel, dela, lead, adle, etc. We can produce all possible permutations ($4! = 4 \times 3 \times 2 \times 1 = 24$) of the word whether they have meaning or not. After manipulation, we can choose which ones have meaning (in some language), but that process is now a semantic one, whereas symbol manipulation is purely syntactic. All signs rely on a certain amount of convention, as all signs have a pragmatic (social) dimension, but symbols are the only signs which require exclusively a social convention, or code, to be understood.

We are used to think of information as pertaining purely to the human realm. In particular, the use of symbolic information, as in our writing system, is thought of as technology used exclusively by humans. Symbols, we have learned, rely on a code, or convention, between symbols and meanings. Such a conventional relation usually specifies rules created by a human community. But it can have a more general definition:

“A code can be defined as a set of rules that establish *a correspondence between two independent worlds*”. The Morse code, for example, connects certain combinations of dots and dashes with the letters of the alphabet. The Highway Code is a liaison between illustrated signals and driving behaviours. A language makes words stand for real objects of the physical World.” [Barbieri, 2003, page 94]

We can thus think of a code as a process that implements correspondence rules between two independent worlds (or classes of objects), by ascribing meaning to arbitrary symbols. Therefore, meaning is not a characteristic of the individual symbols but a convention of the collection of producers and recipients of the encoded information.

Interestingly, we can see such processes in Nature, where the producers and recipients are not human. The prime example is the genetic code, which establishes a correspondence between DNA (the symbolic genes which store information) and proteins, the stuff life on Earth is built of. With very small variations, the genetic code is the same for all life forms. In this sense, we can think of the genetic system and cellular reproduction as a symbolic code whose convention is “accepted” by the collection of all life forms.

Other codes exist in Nature, such as signal transduction from the surface of cells to the genetic system, neural information processing, antigen recognition by antibodies in the immune system, etc. We can also think of animal communication mechanisms, such as the ant pheromone trails, bird signals, etc. Unlike the genetic system, however, most information processes in nature are of an analog rather than digital nature. Throughout this course we will discuss several of these natural codes.

Formalizing Knowledge: Uncovering the Design Principles of Nature²

Once we create symbols, we can also hypothesize relationships among the symbols which we can later check for consistency with what we really observe in the World. By creating relationships among the symbols of things we observe in the World, we are in effect formalizing our knowledge of the World. By formalizing we mean the creation of *rules*, such as verbal arguments and mathematical equations, which define how our symbols relate to one another. In a formalism, the rules that manipulate symbols are independent of their meaning in the sense that they can be calculated mechanically without worrying what symbols stand for.

It is interesting to note that the ability to abstract characteristics of the world from the world itself took thousands of years to be fully established. Even the concept of number, at first was not dissociated from the items being counted. Indeed, several languages (e.g. Japanese) retain vestiges of this process, as different objects are counted with different variations of names for numbers. Physics was the first science to construct precise formal rules of the things in the world. Aristotle (484-322 BC) was the first to relate symbols more explicitly to the external world and to successively clarify the nature of the symbol-world (symbol-matter) relation. “In his *Physics* he proposed that the two main factors which determine an object's speed are its weight and the density of the medium through which it travels. More importantly, he recognized that there could be mathematical rules which could describe the relation between an object's weight, the medium's density and the consequent rate of fall.” [Cariani, 1989, page 52] The rules he proposed to describe this relations were:

1. *For freely falling or freely rising bodies: speed is proportional to the density of the medium.*
2. *In forced motion: speed is directly proportional to the force applied and inversely proportional to the mass of the body moved*

This was the first time that the relationships between observable quantities were hypothesized and used in calculations. Such a formalization of rules as a hypothesis to be tested is what a model is all about. Knowledge is built upon models such as this that sustain our observations of the World.

² This subsection is an excerpt of [Rocha and Schnell, 2005b]

“While these quantities were expressed in terms of numbers, they were still generally regarded as inherent properties of the objects themselves. It was not until Galileo took the interrelationships of the signs themselves as the objects of study that we even see the beginnings of what was to be progressive dissociation of the symbols from the objects represented. Galileo's insight was that the symbols themselves and their interrelations could be studied mathematically quite apart from the relations in the objects that they represented. This process of abstraction was further extended by Newton, who saw that symbols arising from observation [...] are distinct from those involved in representing the physical laws which govern the subsequent motion”. [Cariani, 1989, page 52]

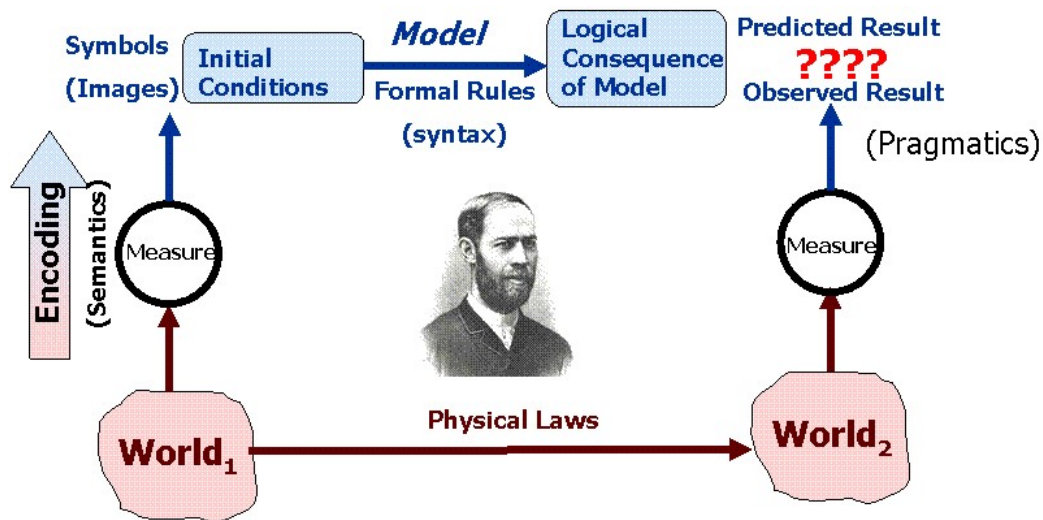


Figure 1: The Modeling Relation between knowledge and reality according to Hertz (adapted from Cariani, 1989)

“In 1894 Heinrich Hertz published *his Principles of Mechanics* which attempted [...] to purge mechanics of metaphysical, mystical, undefined, unmeasured entities such as force and to base the theory explicitly on measurable quantities. Hertz wanted to be as clear, rigorous, and concise as possible, so that implicit, and perhaps unnecessary, concepts could be eliminated from physical theories, [which he thought should be based solely on measurable quantities].” [Cariani, 1989, page 54]. Since the results of measurements are symbols, physical theory should be about building relationships among observationally-derived symbols, that is, it should be about building formal *models*, which Hertz called "images":

“The most direct and in a sense the most important, problem which our conscious knowledge of nature should enable us to solve is the anticipation of future events, so that we may arrange our present affairs in accordance with such anticipation. As a basis for the solution of this problem we always make use of our knowledge of events which have already occurred, obtained by chance observation or by prearranged experiment. In endeavoring thus to draw inferences as to the future from the past, we always adopt the following process. We form for ourselves images or symbols of external objects; and the form which we give them is such that the necessary consequents of the images in thought are always the images of the necessary consequents in nature of the things pictured. In order that this requirement may be satisfied, there must be a certain conformity between nature and our thought. Experience teaches us that the requirement can be satisfied, and hence that

such a conformity does in fact exist. When from our accumulated experiences we have succeeded in deducing images of the desired nature, we can then in a short time develop by means of them, as by means of models, the consequences in the external world which only arise in a comparatively long time, or as a result of our own interposition. We are thus enabled to be in advance of the facts, and to decide as to present affairs in accordance with the insight so obtained. The images which we here speak are of our conceptions of things. With the things themselves they are in conformity in *one* important respect, namely, in satisfying the above mentioned requirement. For our purpose it is not necessary that they should be in conformity with the things in any other respect whatever. As a matter of fact, we do not know, nor do we have any means of knowing, whether our conceptions of things are in conformity with them in any other than the *one* fundamental respect. [Hertz, 1894 pp. 1-2]”

A model is any complete and consistent set of verbal arguments, mathematical equations or computational rules which is thought to correspond to some other entity, its prototype. The prototype can be a physical, biological, social psychological or other conceptual entity.

The etymological roots of the word model lead us to the Latin word “modulus”, which refers to the act of molding, and the Latin word “modus” (a measure) which implies a change of scale in the representation of an entity. The idea of a change of scale, can be interpreted in different ways. As the prototype of a physical, social or natural object, a model represents a change on the scale of abstraction: certain particularities have been removed and simplifications are made to derive a model.

In the natural sciences, models are used as tools for dealing with reality. They are caricatures of the real system specifically built to answer questions about it. By capturing a small number of key elements and leaving out numerous details, models help us to gain a better understanding of reality and the design principles it entails. But it is still surprising how effective mathematical and computational models, these formal caricatures, can aid in the prediction of physical and social reality [Wigner, 1960].

Computational Systems Models vs Thought Experiments ³

“Insofar as the propositions of mathematics are certain they do not refer to reality; and insofar as they refer to reality, they are not certain”. Albert Einstein

Computation is the abstraction and automation of a formal mathematical system, or an axiomatic system. It is defined by the purely syntactic process of mapping symbols to symbols. Such mapping is the basis of the concept of mathematical function, and it is all that computers do. This abstraction requires that all the procedures to manipulate symbols are defined by unambiguous rules that do not depend on physical implementation, space, time, energy considerations or semantic interpretations given to symbols by observers. Formal computation is, by definition, implementation-independent.

Modeling, however, is not entirely a formal process. The Hertzian modeling paradigm (Figure 1) clearly relates formal, computational models to *measurements* of reality against which they must be validated. The measuring process transforms a physical interaction into a symbol – via a measuring device. The measuring process cannot be formalized as it ultimately depends on interacting with a specific (not implementation-independent) portion of reality. We can simulate a measurement process, but for that simulation to be a

³ This section is indebted to many writings of Howard Pattee, including lecture notes and personal communications.

model we will need in turn to relate it to reality via another measurement, and so on indefinitely. This important aspect of modeling is often forgotten in Artificial Life, when the results of simulations are interpreted without access to real world measurements (see Chapter 2).

Likewise, a computer is a physical device that implements a particular abstract computational model as precisely as possible. Modern day computers are so successful because they can implement general-purpose computations almost independently of their specific physics. We do not have to worry about the specific physical architecture of the device as we compute, even though small errors in our computations do occur due to the physical elements of the computing device. In summary, a computation is a process of rewriting symbol strings in a formal system according to a program of rules. The following characteristics are important: (1) Operations and states are syntactic. (2) Symbols follow syntactical rules. (3) Rate of computation is irrelevant. (4) Program determines result, not speed of machine (Physical implementation is irrelevant).

Studying the design principles of life by computational means can be pursued in different ways. While we tend to refer to all of them as simulations, it is important to distinguish approaches that retain empirical validation in the natural, material World, from those that are only theoretically grounded. Only the former are based on Hertz' modeling relation underlying the scientific method. Artificial life, based on the study of *life-as-it-could-be* as proposed by Langton, is circumscribed to theory (see Chapter 2). Its simulations are not scientific models in the sense of Hertz. They are in effect sophisticated computer-aided thought experiments, whose utility is the exploration of the axiomatics, corollaries, and limits of existing theories of life —computational philosophy? Another utility of such computer-aided thought experiments is the generation of bio-inspired machine learning methods, which can also be established by theoretical optimization principles. This engineering discipline does not aim to model life, but rather obtain useful methods to optimize machines and other human and social problems. Interestingly, because the utility of bio-inspired algorithms must ultimately be pragmatically decided in real-World applications, it often provides insights about and challenges to theoretical dispositions more efficiently than the computer-aided thought experiments of artificial life [Floreano & Mattiussi, 2008; Nunes de Castro, 2006].

In systems-theoretic terms, artificial life pertains to the realm of mathematical principles, where isomorphic symbol relabelings are pursued, without measurement or interpretation of experimental observations and validations in the material World [Klir, 2001; Pattee & Raczaszek-Leonardi, 2012]. It is a wholly *systemhood* affair, without direct access to or constraints from *thinghood* [Rosen, 1986]. But as Klir [2001] emphasized, the study of systemhood —without thinghood interpretation or validation — is not (complex) systems *science*, it is mathematics. Indeed, other systems approaches to studying life are more complete in the sense of restricting computational and systems methodologies to *life-as-we-know-it*, thus not getting rid of thinghood altogether. This is the case of the related fields of systems biology, computational biology, bioinformatics, and synthetic biology [Kanehisa, 2000; Kitano, 2002; Villa & Sonis, 2020; Meng & Ellis, 2020], which, like artificial life, also pursue a methodology of holistic synthesis (rather than reductionist analysis). But, unlike artificial life, these disciplines pursue the integration and synthesis of large-scale biochemical, behavioral and ecological information into complex networks and systems models that are validated via measurement of the material World, rather than exclusively from symbolic relabeling within theory.

References

- Aris, R. [1978]. *Mathematical modelling techniques*. London: Pitman. [Reprinted by Dover].
Barbieri, M. [2003]. *The Organic Codes: An Introduction to Semantic Biology*. Cambridge University Press.
Bossel, J. [1994]. *Modeling and simulation*. Wellesley, MA: A K Peters.

- Cariani, P. [1989]. *On the Design of Devices with emergent Semantic Functions*. PhD Dissertation, SUNY Binghamton.
- Hertz, H. [1894]. *Principles of Mechanics*. tr. D. E. Jones and J. T. Walley. New York: Dover 1956.
- Kanehisa, M. [2000]. *Post-genome informatics*. OUP Oxford.
- Kitano, H.. "Systems biology: a brief overview." *Science* **295**.5560 [2002]: 1662-1664
- Klir, G. [2001], *Facets of Systems Science*. Springer.
- Lin, C. C. and Segel, L. A. [1974]. *Mathematics applied to deterministic problems in the natural sciences*. New York: Macmillan Publishing. [Reprinted by SIAM].
- Meng, F., & Ellis, T. [2020]. The second decade of synthetic biology: 2010–2020. *Nature Communications*, **11**[1], 5174.
- 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.
- Polya, G. [1957]. *How to solve it: A new aspect of mathematical method*. Princeton, NJ: Princeton University Press.
- Rocha, L.M. and S. Schnell [2005]. "The Nature of Information". Lecture notes for I101 – Introduction to Informatics. School of Informatics, Indiana University.
- Rocha, L.M. and S. Schnell [2005b]. "Modeling the World". Lecture notes for I101 – Introduction to Informatics. School of Informatics, Indiana University.
- Rosen, R. [1986], "Some Comments on Systems and System Theory". In *Int. Journal of General Systems*. Vol. 13, No.1.
- Saari, D. G. [2001]. *Chaotic elections! A mathematician looks at voting*. American Mathematical Society. Also you can read: "Making sense out of consensus" by Dana MacKenzie, *SIAM News*, Volume 33, Number 8, October 2000.
- Villa, A. & S.T. Sonis [2020]. "System biology." In *Translational Systems Medicine and Oral Disease*, pp. 9-16. Academic Press.
- Wigner, E.P. [1960], "The unreasonable effectiveness of mathematics in the natural sciences". Richard Courant lecture in mathematical sciences delivered at New York University, May 11, 1959. *Comm. Pure Appl. Math.*, **13**: 1-14.

For next lecture read:

- Floreano, D. and C. Mattiussi [2008]. *Bio-Inspired Artificial Intelligence: Theories, Methods, and Technologies*. MIT Press. **Chapter 2**.
- Optional: Nunes de Castro, Leandro [2006]. *Fundamentals of Natural Computing: Basic Concepts, Algorithms, and Applications*. Chapman & Hall. **Chapter 2; Sections 7.3, 8.1, 8.2, and 8.3.10**.
- Optional: Flake [1998], *The Computational Beauty of Life*. MIT Press. **Chapters 10,11, and 14**.

4. Self-Organization and Emergent Complex Behavior

Self-organization is usually understood as the process by which systems of many components tend to reach a particular state, a set of cycling states, or a small volume of their state space (attractor basins), with no external interference. This *attractor behavior* is often recognized at a different level of observation as the spontaneous formation of well-organized structures, patterns, or behaviors, from random initial conditions (emergent behavior). The systems used to study this behavior are referred to as dynamical systems or *state-determined systems*, since every trajectory is perfectly determined by its initial state. Dynamical systems are traditionally studied by continuous variables and sets of discrete-time difference equations (such as the logistic map) or continuous-time differential equations (such as models of the motion of bodies under gravitational forces). However, self-organization is more easily studied computationally with *discrete dynamical systems* (DDS), or automata networks, such as Boolean networks or cellular automata [Rocha, 2022].

The state-determined transition rules of DDS are interpreted as the laws of some physical system [Langton, 1986] where the state of each component depends on the states of its neighbor (or network of interacting) components at the previous time instance. DDS possess a large number of components or variables, and thus very large state spaces. However, when started with random initial conditions (note: not from special initial conditions) they tend to converge, or self-organize, into small sets of attractor states in this space. Attractors may be chaotic in which case the emergent behavior is sensitive to initial conditions. But even chaotic attractors tend to be restricted to small volumes of their state-space (e.g. chaotic in a subset of dimensions of the state-space), therefore we still consider the convergence of a dynamical system into a chaotic basin of attraction to be a form of self-organization.

Since material systems are accurately modeled by dynamical systems, some suggest that it follows from the observed attractor behavior [Wuensche and Lesser, 1992] of these systems that there is a propensity for matter to self-organize (e.g., [Kauffman, 1992]). In this sense, the laws of matter are modeled by the (micro-level) dynamics of state transitions, while its self-organization is studied via the observed (emergent or macro-level) attractor behavior. In general, attractors manifest or *emerge* as global patterns that involve many of components of the dynamical system, and are not easily describable in terms of their state-determined transition rules. For instance, the simple transition rules of the automata in Conway's *Game of Life* cannot describe what the emergent patterns of "blinkers" and "gliders" are. These emergent patterns pertain to a different, complementary level of observation of the same system [Pattee, 1978]. The process of self-organization is often interpreted as the *evolution* of order from random initial conditions. However, notice that this evolution is limited to the specific attractor landscape of a given dynamical system. Unless its parameters are changed (structural perturbation), no dynamical system can escape its own attractor landscape. This limitation will become more apparent when we study the problem of self-replication in Chapter 6.

Life on the Edge of Chaos? The Criticality Hypothesis.

Another interesting aspect of the behavior of dynamical systems concerns the concept of *bifurcation* or *phase transition*. When the parameters of a dynamic system are changed gradually its trajectories and attractors typically change gradually, however, for certain parameter values, sudden changes in the dynamic behavior can occur. It is at this critical point that complicated spatio-temporal organization may occur (e.g. from a steady-state to a limit cycle attractor). Close to bifurcations the system also becomes increasingly more sensitive to parameter and initial condition changes. It is often proposed that bifurcations offer a selection

mechanism [Prigogine, 1985], as a dynamical system may respond very differently to very small changes in their parameters, offering the necessary variation.

However, if the parameter space is divided by many bifurcations, the system becomes increasingly sensitive to initial conditions and small parameter changes; in this sense its behavior becomes chaotic. It has been argued that the richest behavior (greater variation) necessary for evolution and intelligence lies in the transition between order and chaos. Langton [1990, 1992] has shown (for one-dimensional cellular automata) that it is in this range of behavior that dynamical systems can carry the most complicated “computations”. Computation here is used in a loose sense—not as the rate-independent, symbolic manipulation of Turing-machines—meaning that information exchange between elements of these systems is maximized in this range. In other words, Langton showed that the highest correlation among the automata in a lattice network occurs at this stage.

Kauffman [1993,] likewise hypothesized that “living systems exist in the [ordered] regime near the edge of chaos, and natural selection achieves and sustains such a poised state”. This hypothesis is based on Packard’s [1988] work showing that when natural selection algorithms are applied to dynamical systems, with the goal of achieving higher discriminative power, the parameters are changed generally to lead these systems into this transitional area between order and chaos. This idea is very intuitive, since chaotic dynamical systems are extremely sensitive to even small parameter changes, that is, a single perturbation or mutation leads the system into another completely different behavior (sensitive to damage). In contrast, ordered systems are more resilient to damage, and a small parameter change will usually result in a small behavior change which is ideal for smooth adaptation. However, even though very ordered systems can adapt by accumulation of useful successful variations (because damage does not propagate widely), they may not be able ‘step out’ of their particular organization in the presence of novel demands in their environment.

It is here that systems at the edge of chaos were thought to enter the scene; they are not as sensitive to damage as chaotic systems, but still they are more sensitive than fully ordered systems. Thus, most mutations cause only minor structural changes and can accumulate, while a few others may cause major changes in the dynamics enabling a few dramatic changes in behavior. These characteristics of simultaneous mutation buffering (to small changes) and dramatic alteration of behavior (in response to larger changes) is ideal for evolvability [Conrad, 1983, 1990]. However, many of the real gene networks that have been successfully modeled with dynamical systems (e.g. the network of segment polarity genes in *Drosophila melanogaster* [Albert and Othmer, 2003]), appear to be in an ordered regime, being very robust to structural changes [Chaves, Albert and Sontag, 2005; Willadsen&Wiles, 2007; kauffman et al, 2003]. A large study of available experimentally-validated systems biology reveals that most of them are ordered, even after accounting for different ways of measuring dynamical regime [Park et al, 2023]. It appears that evolution favors resilience, i.e. very robust, mostly ordered regimes of self-organization in gene networks – at least the subsystems modeled so far, which could perhaps be biased to highly conserved regulatory pathways. However, the “criticality hypothesis” for biochemical dynamics is very much an active research area [Kauffman, Serra, Shmulevich & Huang, 2023; Balleza et al, 2008]. Indeed, there is evidence that dynamical redundancy in biochemical regulation, whereby biomolecules ignore most of their regulators, facilitates maintenance of near-critical regimes for increased evolvability – even at ranges of connectivity that are expected to yield chaotic dynamics in homogeneous and heterogenous networks [Manicka, Marques-Pita, & Rocha 2022; Costa, Rozum, Marcus & Rocha, 2023], even though how criticality is measured in finite networks is still a matter of debate [Park et al, 2023]. It is also possible that biological subsystems (e.g. relatively smaller pathways or groups of pathways) be ordered, but organisms composed of such systems be critical [Park et al, 2023]. This question may be resolved if and when we get closer to accurate whole organism models.

Complex Self-organization

We have studied several computational systems said to be self-organizing in the sense described above. The discrete logistic equation observes several ranges of ordered behavior according to its parameter r . For $r \leq 3$, the system converges to a single point steady state (independently of its initial value). For $3 \leq r \leq 4$ the system enters a series of bifurcations, meaning that it changes its attractor behavior, first from a steady-state into a two-state limit cycle, and then progressively doubling the number of states in an attractor limit cycle as r increases. Close to $r = 4$, the limit cycle becomes chaotic. That is, in the chaotic range, the slightest change in the initial value, will lead to a completely different trajectory (though similarly chaotic). The system goes from being independent to strongly dependent of initial conditions, though, in each range, the attractor behavior of the equation is the same for random initial conditions. Thus, we can see the logistic equation as self-organizing.

But there is another aspect of the logistic equation that should be understood. In all of its ranges of behavior, from full order to full chaos, the system is (fairly) reversible. That is, I can always obtain a specific initial condition which caused some behavior, by formally running the system backwards. This means the system is deterministic in both temporal directions. Formally, this means the state transition function is invertible. (This is actually only true, if we decide to work on the lower half of its state space, since the logistic equation is a quadratic function, it has always two possible solutions for the previous value of the current state, these values are symmetric about the middle point of its state space). Some, resist calling such of reversible systems self-organizing because they are not sufficiently complex. They reason that if a system is self-organizing, when ran backwards it should be self-disorganizing, that is, it should lead to random initial conditions, or to an incomplete knowledge of possible initial states. Indeed, *complexity* is typically equated with the inability to describe the behavior of a system from the behavior of its components or predecessors. This way, we ought to reserve the term self-organization to those irreversible systems whose behaviors must be evaluated statistically. The logistic map shows “hints” of this backwards self-disorganization, but we can still work out effectively its backwards trajectory to an initial condition by restricting the quadratic solutions to half of its state space.

Random Boolean Networks are much more complicated than this [Gershenson, 2004]. They are completely deterministic since a certain state will always lead to the same next state (state-determinacy), however, we cannot usually know exactly what the predecessor of a current state was. Systems like this are usually studied with statistical tools. Even though the rules that dictate the next state of its components are simple and deterministic, the overall behavior of the system is generally too complicated to predict and statistical analysis has to be performed. For instance, Kauffman [1993] has shown that when $K=2$ (number of inputs to each node), these networks will have on average \sqrt{N} basins of attraction with a length of \sqrt{N} states; if the output of one node is switched to the other boolean value (perturbation), the trajectory returns to that cycle 85% of the time, while on the remaining 15% of the time it will “jump” into a different basin of attraction—denoting (statistical) robustness to perturbation. *Cellular automata* (CA) fall into this same category of deterministic, irreversible, self-organization.

Further Readings and References:

- Albert, R., H.G. Othmer [2003]. “The topology of the regulatory interactions predicts the expression pattern of the segment polarity genes in *Drosophila melanogaster*.” *J Theor Biol.* 223(1):1-18.
- Balleza, E., E. R. Alvarez-Buylla, A. Chaos, S. Kauffman, I. Shmulevich, and M. Aldana. [2008]. “Critical Dynamics in Genetic Regulatory Networks: Examples from Four Kingdoms.” *PLoS One* 3 (6): e2456.
- Costa, F.X.; Rozum, J.C.; Marcus, A.M.; Rocha, L.M. [2023]. "Effective Connectivity and Bias Entropy Improve Prediction of Dynamical Regime in Automata Networks". *Entropy*. 25(2):374. doi: 10.3390/e25020374.

- Chaves, M., R. Albert, and E.D. Sontag [2005]. "Robustness and fragility of Boolean models for genetic regulatory networks." *J Theor Biol.* 235(3):431-49.
- Conrad. M. [1983], *Adaptability*. Plenum Press.
- Conrad, M. [1990], "The geometry of evolutions". In *BioSystems* Vol. 24, pp. 61-81.
- Forrest, S. (Ed.) [1990]. *Emergent Computation*. MIT Press/ North-Holland. Special Issue of Physica D. Vol. 42.
- Gershenson, C. [2004]. "Introduction to random Boolean networks." *arXiv* preprint nlin/0408006.
- Kauffman, S.A. [1993]. *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford University Press.
- Kauffman, Stuart, Carsten Peterson, Björn Samuelsson, and Carl Troein [2003]. "Random Boolean Network Models and the Yeast Transcriptional Network." *Proceedings of the National Academy of Sciences of the United States of America* **100** (25): 14796–14799.
- Kauffman, S.A., R. Serra, I. Shmulevich, S. Huang (Eds.) [2023], *The Principle of Dynamical Criticality. Entropy*. (Special Issue): https://www.mdpi.com/journal/entropy/special_issues/dynamical_criticality
- Klir, George J. [2001]. *Facets of Systems Science*. Springer.
- Langton, C. [1990]. "Computation at the edge of chaos: phase transitions and emergent computation". In Forrest [1990] pp12-37.
- Langton, C. [1992], "Life at the edge of chaos". In *Artificial Life II*. C. Langton (Ed.) Pp 41-91. Addison-Wesley
- Manicka, S., M. Marques-Pita, and L.M. Rocha [2022]. "Effective connectivity determines the critical dynamics of biochemical networks." *Journal of the Royal Society Interface.* **19**(186):20210659. DOI:10.1098/rsif.2021.0659.
- Packard, N. [1988] "Adaptation toward the edge of chaos". In. *Complexity in Biologic Modelling*. S. Kelso and M. Shlesinger.
- Park, K.HH., F.X. Costa, L.M. Rocha, R. Albert, J.C. Rozum [2023]. "Robustness of biomolecular networks suggests functional modules far from the edge of chaos". *PRX Life.* **1**, 023009. DOI: 10.1103/PRXLife.1.023009.
- Pattee, Howard H. [1978]. "The complementarity principle in biological and social structures." In: *Journal of Social and Biological Structures* Vol. 1, pp. 191-200.
- Prigogine, I. [1985]. "New Perspectives on Complexity". In: *Sciences and Praxis of Complexity*. pp.107-118.
- L.M. Rocha [2022]. "On the feasibility of dynamical analysis of network models of biochemical regulation". *Bioinformatics.* **btac360**, DOI: 10.1093/bioinformatics/btac360.
- Willadsen, K. and Wiles, J. [2007] "Robustness and state-space structure of Boolean gene regulatory models". *Journal of Theoretical Biology*, **249** (4), 749-765.
- Wuensche, A., & Lesser, M. [1992]. *The global dynamics of cellular automata: An atlas of basin of attraction fields of one-dimensional cellular automata*. Reading, MA: Addison-Wesley.

5. Reality is Stranger than Fiction

Updated from a [presentation in the “Biocomplexity” discussion](#) section at the 9th *European Conference on Artificial Life*, September 12, 2007 in Lisbon, Portugal

What can Artificial Life do about Advances in Biology?

“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]

From Langton’s original artificial life manifesto, the field was largely expected to free us from the confines of “life-as-we-know-it” and its specific biochemistry. The idea of “life-as-it-could-be” gave us a scientific methodology to consider and study the general principles of life at large. The main assumption of the field was that instead of focusing on the carbon-based, living organization, life could be better explained by synthesizing its “logical forms” from simple machines [Langton, 1989, page 11]—where, “fictional” machines substituted real biochemistry. The expectation was that this “out-of-the-box”, synthetic methodology would deliver a wider scientific understanding of life. We would be able to entertain alternative scenarios for life, challenge the dogmas of biology, and ultimately discover the design principles of life.

Interestingly, during the almost 40 years since the first artificial life workshop, biology witnessed tremendous advances in our understanding of life. True, biology operates at a completely different scale of funding and in a much larger community base than artificial life (the impact factors of key journals in both fields differ by an order of magnitude). But, still, it is from biology, not artificial life, that the strangest and most exciting discoveries and design principles of life arise today. Consider looking at the [September 6, 2007] number of *Nature*, with the quite apropos editorial title “Life as We Know it” [Vol. 449, 1], arguing for a comparative genomics approach, with articles, for instance, moving towards evolutionary principles of gene duplication [Wapinski et al, 2007]. Publications in the [September 2007 issue of] *PLoS. Biol.*, also presented new evidence towards updating or discovering general principles of life: for instance, Venter’s sequencing of his diploid genome, which updates our expectations of differences in chromosome pairs [Levy et al, 2007]), and the Ahituv et al [2007] study that challenges the idea that ultraconserved DNA (across species) must be functional. Since then, many advances, often enabled by big data approaches of computational biology, keep being discovered; for instance, from large-scale comparative genomics, it has been found that retroviral genomic sequences account for 6 to 14% of host genomes—~8% of human DNA is from endogenous retroviruses, which comprises more DNA than the human proteome [Weiss & Stoye, 2013].

It is good to notice that this sort of work is not so much an exception, but has been a signature of research in the biosciences in the last couple of decades. Consider cases such as the discovery of DNA transfer from bacteria to the fly [Dunning Hotopp, 2007], extra-genomic inheritance in *Arabidopsis* [Lolle et al, 2005], or the profound importance of non-coding RNA in life which is a major player in, among other features, patterning [Martello et al, 2007], essential gene regulation [Mattick, 2005], development [Mattick, 2007], epigenetic neural development and modulation [Mehler & Mattick, 2007; Mattick & Mehler, 2008], eukariotic complexity [Taft et al, 2007], etc. Moreover, advances such as these do not seem to be mere epiphenomena of a specific life form. Indeed, they point at important organization principles—as those that artificial life was supposed to provide. When we discover that non-transcribed RNA is involved in extra-genomic inheritance or that most of the evolutionary innovation responsible for differences between

marsupials and placental mammals occurs in non-protein coding DNA [Mikkelsen et al, 2007], some fundamental principles of the living organization are to be re-thought: the simple, generalized genotype-phenotype mappings on which most of artificial life is based on, are just not enough to capture the principles of life as we know it. More intricate genomic structure, and its principles, need to be modeled and theories need to be built to understand life.

One could go on and on about many other advances in biology—*CRISPR* [Ledford, 2017], the *pangenome* [Beavan et al 2024], and even *Horizontal Gene Transfer* involved in “parasitic mind control” [Wilcox, 2023] come to mind, especially as examples of the extended role of DNA and its ability to encode information and facilitate exchanges between very open organisms. We can also point to themes at the forefront of (bio)complexity theory that go largely overlooked in artificial life—though not completely (i.e. [Calabretta et al, 2000; Hintze & Adami, 2007]). Nonetheless, looking at the papers accepted for the main sections of the latest ALIFE and ECAL conferences, it is easy to see that most papers, not only do not discover or even address such issues, but largely trade in biological and computational concepts that have not changed much since the field’s inception (see list of top themes and terms in appendix). Is artificial life trapped in the (evolutionary) biology of forty years ago? Why is reality stranger and more surprising than fiction?

Clearly, there has been very widely successful artificial life research. First and foremost, artificial life has been most successful as a means to study animal behavior, learning and cognition. Certainly, evolutionary robotics and embodied cognition have had an impact in cognitive science. But is artificial life simply a better way to do artificial intelligence? Moreover, one could argue that given the embodied nature of evolutionary robotics, it would seem that it is bound to some kind of material reality, rather than synthesized by constituent “logical forms” as Langton initially suggested.

But what to do about the organization of life itself? Surely the idea of explaining the living organization was behind the origin of the field. For the purposes of this discussion, we must question ourselves why artificial life does not produce more and surprising results about the living organization? Certainly, there is sound research in the field with impact outside of it [e.g. Adami, 2006; Hintze & Adami, 2007]. But even the most successful research in artificial life rarely goes beyond showing that artificial organisms can observe the same behaviors as their real counterparts (i.e. selective pressures, epistasis, etc.). A problem for the field is that as biotechnology gains more and more control of cellular processes, it is reasonable to ask what can one do with artificial organisms that one cannot do with real bacteria? For instance, recent studies of the evolutionary speed towards beneficial mutations were quite effectively done with E-coli [Perfeito et al, 2007], pointing to a much larger rate of beneficial mutations in bacteria than previously thought, and shedding new light on the general principal of clonal interference.

Certainly the community can think of a variety of responses to this lack of new principles of life coming out of research in artificial life—even in theoretical biology. One concept that I venture may need updating in artificial life is its view of the genotype/phenotype relationship. Langton proposed that we generalize this relationship, but this meant that research in the field largely regarded the two as indistinguishable. While this move at first glance seems appropriate to deal with the complexity of genomic-proteomic interaction, it prevents us from studying the specific roles each plays in the living organization. Genotype and phenotype are intertwined in a complex manner, but each operates under different principles that are often overlooked in artificial life. Thus, artificial life rarely approaches issues of genomic structure and regulation, or the co-existence of DNA and RNA as different types of informational carriers. This could well be because artificial life models seem to trade most often on the concept of Mendelian gene than on the molecular biology gene. In other words, artificial life models tend to regard genes solely as mechanisms of generational (vertical)

inheritance, rather than as (informational) mechanisms of ontogenetic (horizontal) development, regulation, maintenance, phenotypic plasticity, and response to environmental change. This way, most artificial life models do not test, or even deal with, possible genomic structure architectures and their impact on development and evolution. This is a big shortcoming in the field since, as we have seen in the last two decades, the molecular biology gene and the genomic structure it implies are behind many essential principles of life—from hypersomatic mutation in vertebrate immunity to speciation.

Additionally, it is most often the case that artificial organisms in artificial life models are designed with many top-down features, rather than emerging out of artificial biochemical machines. For instance, typically the genes of artificial organisms encode pre-defined computer operations. Not only is the encoding pre-defined, but the function of individual genes is also pre-programmed, rather than emergent from some artificial chemistry—what is typically emergent is the behavior of a collection of such “atomic” genes and genotypes.

It is interesting to note that when biologists were looking for the location of genetic information for inheritance, they naturally assumed that it would reside in proteins. They knew of DNA chemically, but its sheer *inertness* deemed it unfit for the job, because of the theoretical expectations set up Schrödinger [1944] that the molecules responsible for genetic inheritance should behave as a “codescript,” i.e. a molecule that could simultaneously function as memory store and engage in catalysis and auto-catalysis. Since catalysis required highly active, dynamical biomolecules, Schrödinger’s disciples, who controlled all early molecular biology, were looking for proteins or something similar. It took some time to realize that relative inertness was really *the* point, from Griffith’s experiment in 1928 to Avery, MacLeod, and McCarty’s in 1944 when DNA was identified as the carrier of genetic information—the implications of which were only fully accepted much later, probably costing Avery a deserved Nobel [Judson, 2003]. Even Avery and team took over a decade to hit on DNA after Griffiths experiment, because they first tried to knock out all cellular constituents that were dynamic. The last thing Schrödinger’s theory allowed him to consider was something so inert as DNA. But of course, inertness is necessary for matter to function as memory (see Chapter 6).

This episode illustrates how reality very often surprises the best scientific expectations of the day. Because Artificial Life defines itself as the study of life-as-it-could-be, it circumscribes itself to what scientists *think* life is, and not on what empirical experiments reveal it to be—and often very much surprise scientists. For instance, the biochemical difference between highly inert memory molecules and highly reactive, functional ones, while often overlooked in artificial life as a design principle, is ultimately the hallmark of life [Rocha and Hordijk, 2005; Brenner, 2012]. Indeed, Venter’s achievement in successfully replicating a living cell with a “*prosthetic genome*” until the original organism’s phenotype is fully re-programmed (see chapter 1), should lead Artificial Life scientists to ponder at least the question of what is it about life’s *design principle* that makes it easier to synthesize a working prosthetic genome than a working “prosthetic” proteome or metabolome? Perhaps, Langton’s view of artificial life being built-up from simple machines, may have clouded the fact that life as we know it is made of biochemical constituents with very different chemical and functional roles: chiefly, DNA (long-term, random-access memory), RNA (short-term memory and symbol processing) and proteins (functional machines). Perhaps more attention should be directed to the “logical forms” of these structural constituents that produce life, before we can tackle “life-as-it-could-be”.

ACKNOWLEDGMENTS

I am extremely grateful to Karola Stotz at Indiana University for exciting conversations that directly lead to this text. I also thank Janet Wiles for performing the *Leximancer* analysis, and Peter Todd for patiently listening as some of the ideas here discussed were germinating.

Further Readings and References

- Ahituv, N., et al. [2007]. “Deletion of Ultraconserved Elements Yields Viable Mice”. *PLoS Biology* Vol. 5, No. 9, e234.
- Beavan, Alan, Maria Rosa Domingo-Sananes, and James O. McInerney. [2024] “Contingency, Repeatability, and Predictability in the Evolution of a Prokaryotic Pangenome.” *Proceedings of the National Academy of Sciences* **121** (1): e2304934120. DOI: 10.1073/pnas.2304934120.
- Brenner, S. [2012]. “Turing centenary: Life’s code script.” *Nature* **482** (7386): 461-461.
- Calabretta, R., Nolfi, S., Parisi, D. and Wagner, G.P. [2000]. “Duplication of modules facilitates the evolution of functional specialization”. *Artificial Life*, 6 , 69-84.
- Hintze, A. and C. Adami [2007]. “Evolution of complex modular biological networks”. arXiv.org:0705.4674.
- Dunning Hotopp, J. C., et al. [2007] "Widespread Lateral Gene Transfer from Intracellular Bacteria to Multicellular Eukaryotes," *Science* doi:10.1126.
- Guimerà, R, Sales-Pardo, M, Amaral, LAN. [2007] “Classes of complex networks defined by role-to-role connectivity profiles”. *Nature Phys.* 3, 63-69.
- Judson, H.F. [2003]. "No Nobel Prize for Whining". *New York Times*, October 20, 2003.
- Ledford, Heidi. [2017] “Five Big Mysteries about CRISPR’s Origins.” *Nature* **541** (7637): 280–82. DOI: 10.1038/541280a.
- Levy, S. et al. [2007] *PLoS Biol.* 5, e254.
- Lolle, S. J., et al [2005]. “Genome-wide non-mendelian inheritance of extra-genomic information in Arabidopsis”, *Nature* 434, 505-509.
- Mikkelsen, T.S. et al [2007]. “Genome of the marsupial *Monodelphis domestica* reveals innovation in non-coding sequences”. *Nature* 447, 167-177. ,
- Martello, G. Et al [2007]. “MicroRNA control of Nodal signalling” *Nature*. doi:10.1038/nature06100;
- Mattick, J.S. [2005] “The functional genomics of noncoding RNA”. *Science*. 309 (5740): 1570-3.
- Mattick JS.[2007]. “A new paradigm for developmental biology”. *J Exp Biol.* ;210(Pt 9):1526-47

- Mattick, JS & MF Mehler [2008] "RNA editing, DNA recoding and the evolution of human cognition". *Trends in Neurosciences*, **31** (5):227-233.
- Mehler MF, Mattick JS. [2007]. "Noncoding RNAs and RNA editing in brain development, functional diversification, and neurological disease". *Physiol Rev.* 87(3):799-823
- Newman, M.E.J. [2006]. "Modularity and community structure in networks". *PNAS.* 103: 8577.
- Perfeito L, L. Fernandes, C. Mota and I. Gordo [2007]. "Adaptive Mutations in Bacteria: High Rate and Small Effects". *Science.* **317** (5839):813 - 815.
- Rocha, L.M. and W. Hordijk [2005]. "Material Representations: From the Genetic Code to the Evolution of Cellular Automata". *Artificial Life.* **11** (1-2):189 - 214
- Schlosser, G. and G. P. Wagner (Eds.) [2004]. *Modularity in Development and Evolution.*
- Schrödinger, Erwin [1944]. *What is Life?* Cambridge University Press.
- Taft RJ, Pheasant M, Mattick JS. [2007]. "The relationship between non-protein-coding DNA and eukaryotic complexity". *Bioessays.* 29(3):288-99.
- Wapinski, I., A. Pfeffer, N. Friedman, and A. Regev [2007]. "Natural history and evolutionary principles of gene duplication in fungi". *Nature* 449, 54-61.1.
- Weiss & Stoye [2013]. "Our Viral Inheritance." *Science.***340** (6134): 820-821.
- Wilcox, Christie. [2023] "Parasitic Worms May Control Minds of Insects with 'Borrowed' Genes." *Science*, DOI: 10.1126/science.adl4678.

6. Von Neumann and Natural Selection

“Turing invented the stored-program computer, and von Neumann showed that the description is separate from the universal constructor. This is not trivial. Physicist Erwin Schrödinger confused the program and the constructor in his 1944 book *What is Life?*, in which he saw chromosomes as “architect’s plan and builder’s craft in one”. This is wrong. The code script contains only a description of the executive function, not the function itself.” [Brenner, 2012]

6.1 Von Neumann’s Theory of Evolvable Self-Reproduction

Von Neumann thought of his logical model of self-reproduction as an answer to the observation that, unlike machines, biological organisms have the ability to self-replicate while seemingly increasing their complexity without limit. Mechanical artefacts are instead produced via more complicated factories (as opposed to self-production) and can only degenerate in their complexity. He was searching for a *threshold of complexity* beyond which machines self-reproduce (with no outside control) while possibly increasing their complexity.

In a posthumous book edited by Arthur Burks, Von Neumann [1966] produced a Cellular Automata implementation of his theory. But the overall theory was introduced in five lectures at the University of Illinois in 1949 [Von Neumann, 1949]. Unlike what Burks’ posthumous book implies (and most people take from the theory), Von Neumann was not looking to show how machines can self-reproduce. In the Illinois lectures, he states that achieving self-replication is trivial, and attainable by non-living structures like crystals which do not grow in complexity. Indeed, the goal was a *logical theory of life* that could explain how organisms locally beat the second law of thermodynamics: a minimum complexity threshold for evolution by natural selection. He clearly stated that the logic he reached (see Figure 1 and explanation below) is not the simplest way to achieve mere self-replication. He was instead after the logical requirements *for evolvable matter*, or machines, via inherited variation.

Von Neumann concluded that the threshold of complexity for evolvable automata (machines) entails a memory-stored description $\Phi(X)$ that can be interpreted by a universal constructor automaton A to produce any automaton X ; when a description of A , $\Phi(A)$, is fed to A itself, then a new copy of A is obtained. However, to avoid a logical paradox of self-reference [Al-Hashimi, 2023], the description, which cannot describe itself, must be both copied (*undecoded* role) and translated (*decoded* role) into the described automaton. This way, in addition to the universal constructor, an automaton B capable of copying any description, $\Phi(X)$, is included in the self-replication scheme. A third automaton C is also included to perform all the manipulation of descriptions necessary—a sort of operating system. To sum it up, the self-replicating system contains the set of automata $(A + B + C)$ and a description $\Phi(A + B + C)$; the description is fed to B which copies it three times (assuming destruction of the original); one of these copies is then fed to A which produces another automaton $(A + B + C)$; the second copy is then handled separately to the new automaton which together with this description is also able to self-reproduce; the third copy is kept so that the self-reproducing capability may be maintained (it is also assumed that A destroys utilized descriptions).

Notice that the description, or memory, is used in two different ways: it is both *translated* and *copied*. In the first role, it controls (or programs) the construction of an automaton by causing a sequence of activities in the machine—Von Neumann named it the *active* role of information. In the second role, the description is simply copied without reference to its meaning—the *passive* role of information. In other words, the *decoded* description controls or programs the construction of a new machine, and the *undecoded* description is copied

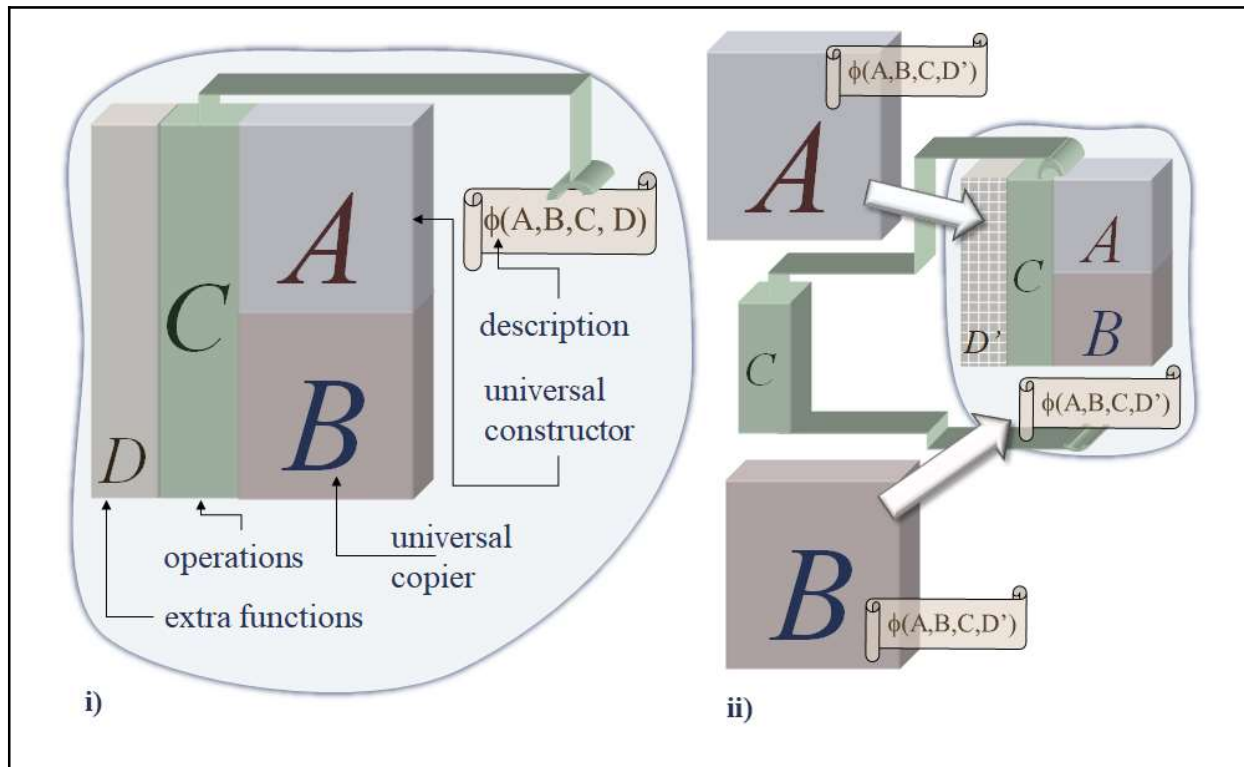


Figure 1: Von Neuman's (threshold of complexity) scheme for evolving machines. i) The four necessary automata and their separate description. ii) A mutation to the description of automaton D' leads to offspring variation in subsequent generation.

separately, passing along its stored information (memory) to the next generation without access to its *meaning*. This parallels the horizontal and vertical transmission of genetic information in biological organisms (protein translation via ribosomes and DNA duplication via polymerases, respectively), which is all the more remarkable since Von Neumann proposed this scheme before the (linear) structure of the DNA molecule was uncovered by Watson, Crick and Franklin [Watson and Crick, 1953]—though after the Avery-MacLeod-McCarty [1944] experiment which identified DNA as the carrier of genetic information.

“The concept of the gene as a symbolic representation of the organism—a code script—is a fundamental feature of the living world and must form the kernel of biological theory.” [Brenner, 2012]

The notion of description-based self-reproduction implies a language. A description must be cast on some *symbol system* while it must also ultimately be implemented by some physical structure (or axiomatic/logical system if considering an exclusively formal treatment) [Rocha & Hordijk, 2005]. When A interprets (or decodes) a description to construct some automaton, a *semantic* code is utilized to map instructions into construction commands to be performed. When B copies a description, only its *syntactic* aspects are replicated—symbol sequences. Now, the language of this semantic code presupposes a set of primitives (e.g. parts and processes) for which the instructions are said to “stand for”. Descriptions are not universal insofar as they refer to these building blocks which cannot be changed without altering the significance of the descriptions. The building blocks ultimately produce the dynamics, behavior, and/or functionality of the overall system. In Biology, the most parsimonious model of protein synthesis regards the genetic code as

instantiating such a language. Genes are one-dimensional, linear descriptions that encode specific parts: amino acids chains which in turn fold into three-dimensional proteins, thus constructing the dynamical machinery of the cell [Pattee, 2022]. In a computational setting, parts are typically logical operations (logical gates), but those must ultimately be implemented as material building blocks in computing hardware, where instructions in the language are translated to specific physical actions [Rocha & Hordijk, 2005].

Von Neumann's active vs. passive modes of information, or the decoded vs undecoded use of genetic descriptions in biology, can also be cast as interpreted vs. uninterpreted material representations [Rocha & Hordijk, 2005], where *interpretation* is understood in the sense of a computer language compiler. However, the term "interpretation" can be confusing and arguably best left for when different alternative meanings for symbols arise at higher levels of organism engagement with an environment, cognition, or social interaction [Pattee, 2021]. For our purposes, what is important is the recognition of the necessity of two distinct roles for descriptions in evolvable self-replication. Per Von Neumann's scheme, and per what we know about biology today [Brenner, 2012; Pattee, 2021, 2022; Al-Hashimi, 2023]: organisms (or evolvable machines) decode descriptions to control material self-construction (horizontally), and also copy (undecoded) descriptions separately and vertically for offspring (though horizontal transmission of undecoded descriptions is also prevalent, as in intercellular RNA transport [Nabariya et al, 202].)

6.2 Open-ended evolution and natural selection

"Biologists ask only three questions of a living organism: how does it work? How is it built? And how did it get that way? They are problems embodied in the classical fields of physiology, embryology and evolution. And at the core of everything are the tapes containing the descriptions to build these special Turing machines." [Brenner, 2012]

Perhaps the most important consequence of *separate* descriptions in Von Neumann's self-reproduction scheme (and Turing's Tape) is its opening the possibility for open-ended evolution [Rocha, 1998; McMullin, 2000]. As Von Neumann [1966] discussed, if the description of the self-reproducing automata is changed (mutated), in a way as to not affect the basic functioning of $(A + B + C)$ then, the new automaton $(A + B + C)'$ will be slightly different from its parent. Von Neumann used a new automaton D to be included in the self-replicating organism, whose function does not disturb the basic performance of $(A + B + C)$; if there is a mutation in the D part of the description, say D' , then the system $(A + B + C + D) + \Phi(A + B + C + D)$ will produce $(A + B + C + D') + \Phi(A + B + C + D')$. Von Neumann [1966, page 86] further proposed that non-trivial self-reproduction should include this "ability to undergo inheritable mutations as well as the ability to make another organism like the original", to distinguish it from "naive" self-reproduction like growing crystals that endlessly reproduce the same form.

Notice that changes in $(A + B + C + D)$ are not heritable, only changes in the description, $\Phi(A + B + C + D)$, are inherited by the automaton's offspring and are thus relevant for evolution. This ability to transmit mutations (vertically) is precisely at the core of the principle of natural selection of modern Darwinism [Richerson, Gavrilets, and de Waal, 2021]. Through variation (mutation) populations of different organisms are produced; the statistical bias these mutations impose on reproduction rates of organisms will create survival differentials (fitness) on the population which define natural selection. In principle, if the language of description is rich enough, an endless variety of organisms can be evolved: *open-ended evolution*.

The evolvability of a self-reproducing system is dependent on the parts used by the semantic code. If the parts are very simple, then the descriptions will have to be very complicated, whereas if the parts possess rich

dynamic properties, the descriptions can be simpler since they will take for granted a lot of the dynamics that otherwise would have to be specified. In the genetic system, genes do not have to specify the functional characteristics of the proteins produced, but simply the string of amino acids that will produce that functionality “for free” [Moreno et al, 1994]. Furthermore, there is a trade-off between programmability and evolvability [Conrad, 1983, 1990] which grants some self-reproducing systems no evolutionary potential whatsoever. When descriptions require high programmability they will be very sensitive to damage. Low programmability grants self-reproducing systems the ability to change without destroying their own organization, though it also reduces the space of possible evolvable configurations [Rocha, 2001].

Considering a varying degree of programmability is crucial to understand the power of the Von Neumann’s theory. He demonstrated the logical feasibility of a universal constructor (A) for the self-replication system using a 29-state cellular automaton [Von Neumann, 1966]. But this specification is best seen as a specific mathematical instance of his general theory of evolvable complexity. More recently, simpler computational implementations of that mathematical proof have been reached [e.g. Pesavento, 1995, see Sipper, 1998]. But Von Neumann famously said that by pursuing implementations of his theory, that is, “by formalizing the problem” he was perhaps “throwing away the baby with the bath water.” Indeed, many people have confused the implementation for the theory, e.g. when suggesting that the overall theory is brittle, and thus not a good model for how living organisms evolve.

The original and current computer implementations are indeed brittle because they are cast in our high-programmability computer languages (e.g. cellular automata), but the qualitative theory put in terms of the A, B, C & D machines outlined above is agnostic as to how one implements it. I side with Pattee [Pattee, & Raczaszek-Leonardi, 2012] and Brenner [2012] in thinking the theory captures quite well (better than any other) what biological cells do, and those are quite robust. In principle, if one adds redundancy to a translation code (like the genetic code) and other components, one can make the artificial implementations of the system much less brittle. Many have argued that the theory should be implemented materially, not logically/computationally to achieve low-programmability and robust emergence of functionality (e.g. Moreno et al [1994], Cariani [1989, 1992]).

The A automaton implementing the decoding/translation function, like any code, is fragile to new semantic assignments. Letters are assigned specific meanings, e.g. "lion" in English or in the genetic code UAA, UAG and UGA signifying the stop codon or AUG the *methionine* amino acid. The semantics of any language are highly affected if those code assignments are tampered with. If by "lion" I mean "chair" English speakers will be confused—and I may even induce someone to be killed by a real lion if I instruct people to sit on them. Similarly, the genetic code is very brittle to changing ribosome/tRNA assignments such as the examples above, which is very likely why we observe almost no variation of the genetic code on earth. However, the genetic code has built-in redundancy to lead to fewer errors in fast translation (only 20 amino acids for 64 possible codon-aminoacid assignments). This adds much robustness in decoding/translation, but not to the ultimate code assignments themselves; if those are tamped with, organisms are likely not viable. In principle, this type of redundancy can be added to artificial implementations of Von Neumann's theory as well [Conrad, 1990]. This would make those less brittle to perturbation in translation, but still as brittle as the genetic code for new code assignments. Indeed, Von Neumann's theory suggests that machines evolved accordingly would have an A automaton with a very conserved code and would only survive if there is built-in redundancy in the code assignments (for reliability under noise).

Turing and Von Neumann were the first to correctly theorize and formalize the required inheritance mechanism behind neo-Darwinian evolution by Natural Selection. This understanding of the most

fundamental design principle of life, puts Turing and Von Neumann on the Parthenon of great thinkers in Biology, alongside Darwin and Mendel [Brenner, 2012]. Our current informational or computational understanding of biochemical machinery allows unprecedented cellular control, as the incredible success of synthetic biology, RNA vaccines, and CRISPR technology demonstrate. Unless a more powerful theory is demonstrated, the dovetailing of computational thinking and biology, inherent in the cybernetics movement of Turing, Von Neumann, Shannon, Wiener and others, emphasizes how (material) control of symbolic information is the hallmark of both computation and biocomplexity.

Further Readings and References:

- Al-Hashimi, H.M. [2024] "Turing, von Neumann, and the computational architecture of biological machines", *Proc. Natl. Acad. Sci. U.S.A.* **120** (25) e2220022120.
- Avery, Oswald T.; Colin M. MacLeod, Maclyn McCarty [1944]. "Studies on the Chemical Nature of the Substance Inducing Transformation of Pneumococcal Types: Induction of Transformation by a Desoxyribonucleic Acid Fraction Isolated from Pneumococcus Type III". *Journal of Experimental Medicine* **79** (2): 137–158.
- Brenner, S. [2012]. "Turing centenary: Life's code script." *Nature* **482** (7386): 461-461.
- Cariani, P. [1989]. *On the Design of Devices with Emergent Semantic Functions*. PhD.Dissertation. SUNY Binghamton.
- Cariani, P. [1992], "Emergence and Artificial Life" In *Artificial Life II*. C. Langton (Ed.). Addison-Wesley. pp. 775-797.
- Conrad. M. [1983], *Adaptability*. Plenum Press.
- Conrad, M. [1990], "The geometry of evolutions". In *BioSystems* Vol. 24, pp. 61-81.
- McMullin, B. [2000]. "John von Neumann and the Evolutionary Growth of Complexity: Looking Backwards, Looking Forwards". *Artificial Life* 6(4):347-361.
- 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.
- Nabariya, D. K., Heinz, A., Derksen, S., & Krauß, S. [2022]. "Intracellular and intercellular transport of RNA organelles in CXG repeat disorders: The strength of weak ties". *Frontiers in molecular biosciences*, **9**, 1000932.
- Pattee, H.H. and J. Raczaszek-Leonardi.[2012]. *Laws, language and life: Howard Pattee's classic papers on the physics of symbols with contemporary commentary*. Vol. 7. Springer.
- Pattee, H. H. [2021] "Symbol Grounding Precedes Interpretation." *Biosemiotics* **14** (3): 561-568.
- Pattee, H.H [2022]. "The Primary Biosemiosis: Symbol Sequence Grounding by Folding". In *Open Semiotics*, Biglari, A. (Ed.) Paris: L'Harmattan.
- Pesavento, U. [1995] An implementation of von Neumann's self-reproducing machine. *Artificial Life* 2(4):337-354.
- Richerson, P.J., S. Gavrilets, and Frans B. M. de Waal [2021] "Modern theories of human evolution foreshadowed by Darwin's Descent of Man". *Science* **372**, eaba3776.
- Rocha, L.M. [1998]. "Selected Self-Organization and the Semiotics of Evolutionary Systems". In: *Evolutionary Systems: The Biological and Epistemological Perspectives on Selection and Self- Organization*, . S. Salthe, G. Van de Vijver, and M. Delpo (eds.). Kluwer, pp. 341-358.[2]
- Rocha, L.M. [2001]. "Evolution with material symbol systems". *Biosystems*. 60, pp. 95-121.
- Rocha, Luis M. and W. Hordijk [2005]. "Material Representations: From the Genetic Code to the Evolution of Cellular Automata". *Artificial Life*. **11** (1-2), pp. 189 - 214
- Sipper, M. [1998]. "Fifty Years of Research on Self-replication: an Overview". *Artificial Life*, 4 (3), pp. 237-257.
- Von Neumann, J. [1949]. "Theory and organization of complicated automata." 5 lectures at University of Illinois. Included in [Von Neumann, 1966].
- von Neumann, John [1966]. *The Theory of Self-Reproducing Automata*. Arthur Burks (Ed.) University of Illinois Press.
- Watson Jd, Crick Fh [1953]. "Molecular structure of nucleic acids; a structure for deoxyribose nucleic acid". *Nature* **171** (4356): 737–8.

7. Modeling Evolutionary Systems

“How does evolution produce increasingly fit organisms in environments which are highly uncertain for individual organisms? How does an organism use its experience to modify its behavior in beneficial ways (i.e. how does it learn or ‘adapt under sensory guidance’)? How can computers be programmed so that problem-solving capabilities are built up by specifying ‘*what* is to be done’ rather than ‘*how* to do it’?” [Holland, 1975, page 1]

These were some of the questions concerning John Holland when he thought of Genetic Algorithms (GA’s) in the 1960’s. All these questions were shown to be reducible to a problem of optimizing multi-parameter functions. Nature’s “problem” (that was addressed formally and conceptually by Von Neumann’s threshold of complexity as described in Chapter 6) is to create organisms that reproduce more (are more fit) in a particular environment: the environment-organism coupling dictates the selective pressures, and the solutions to these pressures are organisms themselves. In the language of optimization, the solutions to a particular problem (say, an engineering problem), will be selected according to how well they solve that problem. GA’s are inspired by natural selection as the solutions to our problem are not algebraically calculated, but rather found by a population of solution alternatives which is altered in each time step of the algorithm in order to increase the probability of having better solutions in the next reproduced population. In other words, GA’s, or other Evolutionary Strategies (ES) such as Evolutionary Programming (EP), explore the multi-parameter space of solution alternatives for a particular problem, by means of a population of encoded strings, or descriptions (standing for alternatives) which undergo variation (crossover and mutation) and are reproduced in a way as to lead the population to ever more promising regions of this search space (selection) [Goldeberg, 1989; Mitchell, 1999; De Jong, 2006].

7.1 Evolutionary Strategies and Self-Organization

The underlying idea of computational ES is the *separation* of solutions for a particular problem (e.g. a machine) from *descriptions* of those solutions (memory). GA’s work on the descriptions and not on the solutions themselves, that is, variation is applied to descriptions, while the respective solutions are evaluated, and the whole (description-solution) selected according to this evaluation. Such machine/description separation follows von Neumann’s self-reproducing scheme (see chapter 6) which is able to increase the complexity of the (organization of) machines described. In this sense, evolutionary algorithms are a specification (for a specific finite problem) of the general system of open-ended complexity setup by Von Neumann—Klir [2001] would refer to this process as an *interpretation* of the general system.

The organizing principle underlying GAs is very different from the principles of self-organization, e.g. in the sense of a boolean network or cellular automata (see chapter 4). Even though solutions are obtained from the interaction of a population of elements, and in this sense following the general rules usually observed by computationally emergent systems (e.g. Langton [1988], Mitchell [1992]), they do not *self*-organize since they rely on the external *selective* pressures of some environment (in ES, defined by an explicit or implicit fitness function). The order so attained is not a result of the internal dynamics of a collection of interacting elements, but is instead dictated by the *external* selection criteria. In this sense, ES follow an organizing scheme that is driven by external selection of encoded symbolic descriptions (a “Turing tape” [Brenner, 2012]). It is perhaps useful to think that ES are modeling the most fundamental design principle of biological systems: natural selection. While self-organizing systems model the dynamical characteristics of matter, ES model the existence of external, selective pressures on populations of symbolic descriptions of (encoded) machines. While self-organization models material dynamics, ES model the selection, variation, and

inheritance of information (encoded in stable, effectively inert, memory structures [Rocha and Hordijk, 2005]) about dynamics.

7.2 Development and morphogenesis: self-organization and selection come together

Since the original introduction of GA's, many subsequent advancements focused on the inclusion of a developmental stage, or intermediate layers between genotype and phenotype; in other words, the creation of some artificial ontogeny [Buckley, 2008]. Since GA depend on a code between symbolic descriptions and machines, developmental regulation can act on either side of the code. Thus, one possibility is to encode post-translation rules that allow decoded products to self-organize to produce a phenotype, rather than the direct encoding of the phenotype itself. Another idea is the introduction of pre-translation (post-transcription) genomic regulation. For instance, post-translation rules may use L-System grammars which dictate production system programs [Wilson, 1988] leading to some phenotype. The most important advantage of this intermediate stage, as initially explored by Kitano [1990], Gruau [1993], Belew [1992] and others, is the ability to code for much larger structures than a direct encoding allows. In practical terms, they have solved some of the scalability problems of encoding (e.g.) neural networks in GA's, by reducing the search space dramatically [Stanley et al, 2019].

The first approaches to model self-organizing developmental processes used solely context-free, state-determined, L-System grammars, compromising epistasis (or mutual, non-linear, influence of genetic descriptions amongst each other) in the simulation of self-organizing development. Dellaert and Beer [1994] and Kitano [1994], for instance, used Boolean networks to simulate genetic epistasis and self-organization. In other words, the GA encodes developmental rules which construct Boolean networks whose nodes stand for aspects of the phenotypes we wish to evolve on some physical simulation. In Dellaert and Beer's model, the nodes stand for cell mitosis and other characteristics. This way, the solutions of the GA are self-organizing systems whose attractor behavior dictates pre-defined phenotypic traits. These approaches in effect offer an emergent morphology, that is, they encode rules which will themselves self-organize into some phenotype (instead of strict programming of morphology). The indirect encoding further allows the search to occur in a reduced space, amplified through development. An interesting side effect has to do with the appearance of modularity traits on the evolved phenotypes [Wagner, 1995]. Subsequent developments paid even more attention to the contextual regulation that indirect encodings afford to the search [Rocha 1995, 1997, 2001].

More recently, given our expanded view of genomics, pre-translation intermediate layers between genotype and phenotype have been explored, such as transcription regulation [Reil, 1999; Hallinan & Wiles, 2004] and RNA Editing [Rocha et al, 2006]. The inclusion of more sophisticated regulation of genetic information prior to translation, while not necessarily including a self-organizing component, allows us to model a much more realistic genotype/phenotype/environment interaction. Instead of genotypes used exclusively for Mendelian inheritance (see chapter 5) of (directly encoded) phenotypic traits, ES with genotype regulation allow us to model the contextual, plastic development of phenotypes we have come to understand via modern Genomics—thus also learning additional design principles for bio-inspired computation [Huang et al, 2007].

The most important aspect of GA's with emergent morphologies is the utilization in the same model of an external selection engine (the GA) coupled to a particular self-organizing dynamics (e.g. Boolean networks) standing for some materiality. Such schemes bring together, computationally, the two most important aspects of evolutionary systems: self-organization and selection. These models belong to a category of self-

organization referred to as *Selected Self-Organization* which is based on symbolic memory [Rocha, 1996, 1997, 1998]. Selected self-organization is also possible in purely dynamical, autocatalytic structures, though its evolutionary potential is much more limited than evolution via symbolic encoding in linear memory structures like DNA (externalized Turing tapes) [Rocha, 2001][Vasas, 2010]. The reason lies in Von Neumann's requirements for evolvable complexity (see chapter 6). The introduction of symbolic descriptions allows a much more powerful form of communication: dynamical structures are constructed from static descriptions and do not have to reproduce through some dynamically complicated, and limited process of self-inspection for self-reproduction [Rocha, 2001]. In other words, separate descriptions (stored in effectively inert memory structures like DNA) can be used to reliably construct any dynamical structure (from a pool of building blocks like amino acids) in an open-ended manner, while self-inspection relies on only those structures that happen to be able to make copies of themselves. As an example, a non-genetic protein-based life form, would have to rely only on those proteins that could make direct copies of themselves, which is much difficult to attain and severely limits what is possible to evolve [Rocha, 2001].

7.3 Multiscale complexity of evolutionary systems

Throughout this course emphasis was put on identifying the most important tools utilized in the fields of Artificial Life, Bio-Inspired Computing, and Evolutionary (Complex) Systems. We started with *self-organizing systems*, exemplified with the logistics equation, random boolean networks, cellular automata (e.g. Conway's game of Life), and all characterized in terms of *dynamical systems theory*. Later, with the *von Neumann's evolvable complexity scheme*, I argued that state-determined (purely dynamic) systems are not able to offer *open-ended evolution*, that is, to increase their complexity with genuine emergence of new functionalities. Dynamic systems are restricted to the complexity of their attractor landscape.

For this purpose, systems inspired by von Neumann's scheme, which demand the separation between the description of a machine from the machine itself, and therefore introduce the concept of *symbolic memory* and external selection, were introduced [Brenner, 2012; Pattee, 2022]. Such systems offer a model of the mechanisms utilized by natural selection, and are accordingly known as *evolutionary systems* (or evolutionary strategies) — e.g. genetic algorithms, evolutionary programming, and the wider class computational models of evolution. We can also refer to the mechanisms utilized to model the kind of evolution that natural selection offers as *memory-based selective strategies*: selection acting on memory elements in order to change the dynamic structure they encode.

I further emphasized hybrid systems which try to model both the self-organizing and selective mechanisms of biological systems, and can therefore offer a more complete understanding of evolutionary systems. These memory-based, selected self-organizing systems, follow a view of evolutionary systems based on the *semiotic closure* of genotype-phenotype maps [Pattee & Raczaszek-Leonardi, 2012], which are rather open to control [Rocha, 2000] and prone to be organized in multiple scales of interaction (multi-level complex networks) [Pattee, 1973; Wilson, 2020]. For practical applications, I showed approaches aiming at the introduction of non-deterministic, self-organizing, developmental steps between genotype and phenotype such as the evolution of boolean/neural networks encoded through L-System rules in a genetic algorithm. Also discussed were models capable of emergent computation by coupling genetic algorithms to cellular automata in order to have the latter solve non-trivial tasks, and the dynamics of adaptive immunity that harnesses a "private" natural selection of lymphocytes in vertebrates.

The understanding of the relative importance the two basic categories of organization in artificial systems introduces a very powerful way to study the relative importance of self-organization and natural selection

in biological systems themselves. In other words, by creating different forms of *life-as-it-could-be* with different degrees of both these categories, we may shed some light on the credit assignment problem of biology: how much of evolution is a result of natural selection and how much is a result of the self-organizing characteristics of its specific materiality.

I was able to introduce many of the usual applications of Artificial Life and Bio-Inspired Computing, from bugs and boids, to evolutionary robots, social evolution, and the multi-level complexity of adaptive immunity. Each of these applications can be a universe of investigation in itself, so emphasis was instead put on the basic categories of organization and their respective simulation tools referred above. In one way or another, all of these applications utilize in different degrees such tools described throughout the course. For instance, evolutionary robots may use a genetic algorithm to evolve a (self-organizing) neural network for its control system allowing it to solve some maze. To the extent that its control system was evolved and uses self-organizing mechanisms, we can say that such control system was subjected to a memory based selective type of self-organization. Naturally, the robot itself (its moving parts and sensors) were not evolved but engineered; the complete evolution of a robot through self-organization and selection represents probably the most ambitious long-term goal of Artificial Life, showing us how far behind we still are from getting there.

Further Readings and References:

- Belew, R.K. [1993]. "Interposing an Ontogenic Model Between Genetic Algorithms and Neural Networks." In: *Advances in neural information processing (NIPS5)*. J. Cowan (Ed.). Morgan Kaufmann.
- Brenner, S. [2012]. "Turing centenary: Life's code script." *Nature* **482** (7386): 461-461.
- Buckley, W.R. [2008] "Computational ontogeny." *Biological Theory* **3**(1): 3-6.
- De Jong, K.A. [2006]. *Evolutionary Computation: A Unified Approach*. MIT Press.
- Dellaert, F. and R.D. Beer [1994]. "Toward an evolvable model of development for autonomous agent synthesis." In: *Artificial Life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*. R. Brooks and P. Maes (Eds.). MIT Press.
- Goldberg, D. E. [1989]. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley.
- Gruau, Frédéric [1993]. "Genetic Synthesis of Modular Neural Networks." In: *Proceedings of the fifth international conference on Genetic Algorithms*. S. Forrest. Morgan Kaufman, pp 318-325.
- Hallinan, J. And J Wiles [2004]. "Asynchronous Dynamics of an Artificial Genetic Regulatory Network". *9th International Conference on Artificial Life*. MIT Press.
- Holland, John H. [1975]. *Adaptation in Natural and Artificial Systems*. University of Michigan Press.
- C. Huang, J. Kaur, A. Maguitman, L.M. Rocha [2007]. "Agent-Based Model of Genotype Editing". *Evolutionary Computation*, 15(3).
- Kitano, H. [1990]. "Designing Networks using Genetic Algorithms with Graph Generation System." In: *Complex Systems* Vol. 4, pp 461-476.
- Kitano, Hiroaki [1994]. "Evolution of Metabolism for Morphogenesis." In: *Artificial Life IV: proceedings of the fourth international workshop on the synthesis and simulation of living systems*. R. Brooks and P. Maes (Eds.). MIT Press.
- Klir, G. [2001], *Facets of Systems Science*. Springer.
- Mitchell, M. [1992]. "Genetic algorithms". In *Lectures in Complex Systems*. L. Nadel and D. Stein (Eds.). SFI Studies in the Science of Complexity Vol. V, Addison-Wesley. pp 3-87.
- Mitchell, M. [1999]. *An Introduction to Genetic Algorithms*. MIT Press.
- Pattee, H. H. (1973). *Hierarchy theory: The Challenge of Complex Systems*. Braziller.
- 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.
- Pattee, H.H [2022]. "The Primary Biosemiosis: Symbol Sequence Grounding by Folding". In *Open Semiotics*, Biglari, A. (Ed.) Paris: L'Harmattan.


- Reil, T. 1999. "Dynamics of Gene Expression in an Artificial Genome - Implications for Biological and Artificial Ontogeny". In *Proceedings of the 5th European Conference on Advances in Artificial Life* (September 13 - 17, 1999). D. Floreano, J. Nicoud, and F. Mondada, Eds. Lecture Notes In Computer Science, vol. 1674. Springer-Verlag, London, 457-466
- Rocha, Luis M. [1995]. "Contextual Genetic Algorithms: Evolving Developmental Rules." In: *Advances in Artificial Life*. F. Moran, A. Moreno, J.J. Merelo, and P. Chacon (Eds.). Series: Lecture Notes in Artificial Intelligence, Springer-Verlag, pp. 368-382.
- Rocha, Luis M. [1996]. "Eigenbehavior and symbols." In: *Systems Research* Vol. 12, No. 3, pp. 371-384.
- Rocha, Luis M. [1997]. *Evidence Sets and Contextual Genetic Algorithms: Exploring Uncertainty, Context, and Embodiment in Cognitive and Biological Systems*. PhD. Dissertation. SUNY Binghamton.
- Rocha, Luis M. [1998]. "Selected Self-Organization and the Semiotics of Evolutionary Systems." In: *Evolutionary Systems: The Biological and Epistemological Perspectives on Selection and Self-Organization*. S. Salthe, G. Van de Vijver, and M. Delpos (eds.). Kluwer Academic Publishers, pp. 341-358.
- Rocha, Luis M. [2000]. "Syntactic autonomy, cellular automata, and RNA editing: or why self-organization needs symbols to evolve and how it might evolve them". In: *Closure: Emergent Organizations and Their Dynamics*. Chandler J.L.R. and G, Van de Vijver (Eds.) *Annals of the New York Academy of Sciences*. Vol. 901, pp 207-223.
- Rocha, Luis M. [2001]. "Evolution with Material Symbol Systems". *Biosystems*. Vol. 60, pp. 95-121.
- Rocha, Luis M. and W. Hordijk [2005]. "Material Representations: From the Genetic Code to the Evolution of Cellular Automata". *Artificial Life*. **11** (1-2), pp. 189 - 214
- Rocha, L.M., A. Maguitman, C. Huang, J. Kaur, and S. Narayanan. [2006]. "An Evolutionary Model of Genotype Editing". In: *Artificial Life 10: Tenth International Conference on the Simulation and Synthesis of Living Systems*. L.M.Rocha, L. Yaeger, M. Bedau, D. Floreano, R. Goldstone, and A. Vespignani (Eds.). MIT Press, pp. 105-111.
- Stanley, K. O., Clune, J., Lehman, J., & Miikkulainen, R. (2019). Designing neural networks through neuroevolution. *Nature Machine Intelligence*, **1**(1), 24-35.
- Vasas, Vera, Eors Szathmary, and Mauro Santos [2010]. "Lack of evolvability in self-sustaining autocatalytic networks: A constraint on the metabolism-first path to the origin of life. ". *Proceedings of the National Academy of Sciences of the United States of America*: 0912628107.
- Wagner, Gunter [1995] "Adaptation and the modular design of organisms". In: *Advances in Artificial Life*. F. Moran, A. Moreno, J.J. Merelo, and P. Chacon (Eds.). Series: Lecture Notes in Artificial Intelligence, Springer-Verlag, pp. 317-328.
- Wilson, D. S. [2020]. *This view of life: Completing the Darwinian revolution*. Vintage.

APPENDIX:

Top themes extracted from all abstracts accepted to ECAL 2007, produced the Leximancer (courtesy of Janet Wiles)

Concept	Absolute Count	Relative Count
<u>model</u>	119	100%
<u>system</u>	95	79.8%
<u>evolution</u>	92	77.3%
<u>results</u>	64	53.7%
<u>environment</u>	58	48.7%
<u>behavior</u>	57	47.8%
<u>present</u>	57	47.8%
<u>network</u>	56	47%
<u>robot</u>	55	46.2%
<u>agents</u>	52	43.6%
<u>simulation</u>	50	42%
<u>process</u>	49	41.1%
<u>simple</u>	46	38.6%
<u>show</u>	43	36.1%
<u>mechanism</u>	41	34.4%
<u>complex</u>	41	34.4%
<u>dynamics</u>	39	32.7%
<u>artificial</u>	39	32.7%
<u>problem</u>	39	32.7%
<u>based</u>	36	30.2%
<u>learning</u>	36	30.2%
<u>approach</u>	33	27.7%
<u>population</u>	31	26%
<u>study</u>	31	26%
<u>genetic</u>	30	25.2%
<u>individual</u>	30	25.2%
<u>neural</u>	28	23.5%
<u>selection</u>	28	23.5%
<u>organisms</u>	27	22.6%
<u>method</u>	26	21.8%
<u>conditions</u>	26	21.8%
<u>level</u>	26	21.8%
<u>information</u>	25	21%
<u>social</u>	25	21%

Top Themes produced from Leximancer set at 65% coverage themes (courtesy of Janet Wiles)



Top co-occurring (stemmed) word pairs in abstracts

neural--network
chang--environ
artifici--life
simul--result
autonom--robot
evolutionari--algorithm
evolutionari--robot
comput--simul
genet--algorithm
robot--mobil
cellular--automata
interact--between
artifici--chemistri

agent--adapt
pressur--select
neural--CONTROL