# Introduction: Living Systems, Causality, and Information

While it is clear that life and living systems are fascinating and complex, the reader may still ask why we talk about living systems in the context of causality and information. You now have a text titled *Algorithmic Information Dynamics – A Computational Approach to Causality and Living Systems* in front of you. This may appear an unlikely constellation of terms. Computation is not something you readily associate with life or living things. Causality and information are also usually treated in very different domains of science and research. Life and causality appear not to be that close either. Now, let us disentangle some of the reasoning behind our assembling what appear to be an ill-assorted melange of terms.

First, some words on living systems. Fundamentally, from a physics standpoint it is not evident how we are to understand life and the emergence of autonomous living systems on the basis of fundamental equations such as those of Schrödinger and Maxwell. To formulate the problem in these terms may not yield a productive research program despite the fact that living systems must obey the fundamental laws of physics. Asking how to construct or generate models of living systems from these equations does not appear to be a pragmatic way forward considering the currently impoverished state of our fundamental understanding of living systems. Historically, the quest for an understanding of living systems, in contrast to what can be referred to as dead systems, has fuelled development of different scientific and non-scientific views and ontologies. Scientifically, we can broadly discern a path originating in Aristotle's teleological biology, including his analysis of four different types of causes, leading to different types of explanations of phenomena in the world. Aristotle had realised that biological systems appeared to be fundamentally different from rocks and stones. The next conceptual shift came with the publication in 1859 of Darwin's *On the Origin of Species*. It was Darwin's genius was natural selection as a sufficient driver and producer of biological diversity, a claim that has since been supported by ample empirical evidence. An element of randomness, altering a code, resulting in modified phenotypes (organisms) is in this sense sufficient to generate diversity. Now, at this juncture one can ask several questions: about the source of randomness, about the kind of code involved. The next chapter in this history – addressing the issue of the code – was written in the middle of the last century, when the structure of DNA was deciphered and proved to be a vehicle to store and transfer genetic information, which turned out to be fundamentally a discrete computational representation. Following the complete sequencing of the human genome and the realisation of the sheer complexity in genome organisation and

https://doi.org/10.1017/9781108596619.002 Published online by Cambridge University Press

magnitude of the various kinds of biomolecules, we saw the renaissance of systems biology. Making the transition between reading the discrete genomic code and assessing the impact on the phenotype, the Darwinian organism, proved to be an exceedingly complex business. Systems biology can therefore be succinctly characterised as an integrative approach, albeit one grounded in biomolecules and data, that tries to grapple with the complexity of living systems in this specific sense. The problem can be deconstructed into several steps, each tremendously challenging. First, how does a code produce a diversity of different cells? Next, how does the collective behaviour of cells eventuate in tissues and organs, and ultimately in a functional living system capable of surviving in and acting on its environment? In parallel with these selected historical stepping stones, there has been an unprecedented development of algorithms, architecture, and hardware since 1936, when Turing and others offered a pioneering conception of computation. Of particular relevance for the understanding of living systems is the mathematical exploration of necessary and sufficient conditions for self-reproducing machines by John von Neumann in the 1950s using cellular automata. Fundamentally, a living system has to somehow sustain and rebuild itself over time. The genius of von Neumann was first to ask how this happened, and secondly to develop a formalism with which to address the question in a precise manner. Even though von Neumann was a leading expert in quantum mechanics, it's important to realise that he understood that breathing life, as it were, into such equations to produce a self-reproducing system, was neither practical nor conceptually attractive. Instead, his approach opened up to ask: What is the minimal program sufficient to generate a self-reproducing system? Subsequently, there have been numerous research programs targeting topics such as complex systems, self-organisation, and non-linear dynamics – in computer science, chemistry, and mathematics. Thus we can safely assume that notions such as information, complex systems, and causality cannot be ignored in the analysis of living systems, and indeed they should not be sealed off from such notions.

Let us be careful about the kinds of questions we can ask about living systems, and precise in specifying the particular questions we wish to address. This sets the stage for understanding the architecture of the current book. In principle, we can ask four different types of questions:

A: Historical: What is the origin of life? This can be elaborated on in a number of ways. When did it originate? How we define life affects this temporal question, as well as the question of how life developed from simpler components. Did life originate on earth or in space? Addressing the order of events and the primacy of different types of molecules is still fertile ground for arguments and experiments. Is our world an RNA world or not?

B: Temporal progression of species: How did life develop? What are the drivers of evolution beyond or behind natural evolution? Current issues include elucidating molecular mechanisms as well as understanding species dynamics over space and time in response to different environmental conditions.

C: Causation and the mechanisms by which information is transferred across generations: With the discovery of DNA, we now have a blueprint for how this works. The

https://doi.org/10.1017/9781108596619.002 Published online by Cambridge University Press

problem is solved in the sense that DNA is the carrier, whereas massive amounts of work has been carried out since the mid-twentieth century on exactly how this works in biochemical detail. A hot topic generating a lot of interest is the notion of transgenerational transfer as a consequence of acquired epigenetic changes.

D: Dynamical, i.e., the stability and sustainability of living systems: This amounts to asking why living systems are not a molecular house of cards that falls apart given the second law of thermodynamics. What are the mechanisms that are sufficient to generate systems that are alive, robust, self-reproducing, and flexible enough to learn and adapt from their local environments?

First, we would like to underline that these are very different types of questions, and that each can be addressed using different techniques originating in different professional fields. Secondly, it is interesting that Schrödinger effectively addressed C and D in his classic book *What Is Life?* published in 1944. As a physicist, he speculated, prior to the discovery of DNA, that for problem C, the answer had to be some kind of repetitive crystal structure transmitting the information (the code). A decade later, DNA was established as just such a structure, which interestingly turned out to be discrete. Furthermore, problem D is conceptually similar to an earlier foundational problem in physics. Physicists had been intrigued by how the atom could be a stable structure, a challenge which required the development of quantum mechanics before it could be addressed. In brief, why does the periodic table look the way it does? Who can interact with whom, and what are the fundamental rules determining the reactivity and stability of the different elements? Thus, it is rather natural to ask a similarly fundamental question – as Schrödinger did – about what accounts for the stability of biomolecules defining molecular machines and, by extension, for the persistence of living systems – functional metabolic circuits, cells, tissues, organs, and eventually organisms surviving in an environment.

In our book, our initial entry point is essentially problem D, viewing living systems as complex systems. We search for a conceptual and practical approach to assess what they are, how they work, and how such systems can extract regularities from their environment. Extracting regularities is a consequence of the fact that the world is not random; there are causal regularities out there. Without such regularities it is difficult to conceive of the existence of living systems, let alone Darwinian evolution. Indeed, the existence of causal regularities in the world appears necessary for the existence of living systems, as how would they survive otherwise? Furthermore, such systems benefit from having the ability to detect regularities and exploit them in their actions. Naturally, since living systems are part of such a world of regularities, their very construction reflects causal dependencies. This outlook – which motivates our work – therefore leads us to consider topics such as causality, complex systems, and information from a fundamental point of view, and to specifically address how they are intertwined, this intertwining being precisely the challenge that the existence of living systems poses.

Let us briefly digress and consider the notions of complex systems, networks, and causality, which will eventually set the stage for using the concept of information as an organising principle upon which to base our analysis. On the one hand we have learned,

https://doi.org/10.1017/9781108596619.002 Published online by Cambridge University Press

ever since the pioneering studies of Poincaré, that non-linear dynamical systems can exhibit chaotic behaviour as well as produce regular repeating patterns, which, however, are not predictable from a microscopic point of view. Poincaré was also interested in, or rather challenged by the Swedish King to tackle, the question of stability. Given Newtonian dynamics, are the planetary orbits stable or not? The famous mathematician Mittag-Leffler headed the prize committee and this was the inception of the famous Mittag-Leffler Institute for Mathematics in Stockholm. In response to the challenge and building on his ongoing work, Poincaré invented the field of dynamical systems. In the 100 years since, stunning progress has been made in understanding phenomena such as sensitivity to initial conditions, bifurcations, and attractors using some of Poincaré's tools. The key point for our purposes is that systems, which from an equation point of view appear simple, can indeed exhibit very complicated chaotic dynamical behaviour. The flip side of the coin is that very complicated systems of equations can effectively display low-dimensional dynamical behaviour, which can be captured in reduced, simpler mathematical descriptions. These yin-yang, complex-simple facets of non-linear systems have been extremely useful stepping stones toward further exploration. Now we have a mathematical language to describe and analyse such dual phenomena in various complex systems. For example, modern machine learning techniques such as deep convolutional networks detect low-dimensional representation in data, akin to techniques such as principal component analysis (PCA), and matrix decomposition algorithms more generally. Exploration of the sensitivity and complexity of chemical systems has been pioneered by Prigogine in his Nobel Prize work (1977) on non-equilibrium systems, offering one example of how to ground complex non-linear experimental systems using such language as non-linear dynamics. These systems are, as a rule, simple in the sense that their dimensionality is limited. The number of equations is either rather small, say one or two orders of magnitude (non-linear dynamics in applied mathematics), or alternatively, the systems are homogeneous (e.g., spin glass models in physics), thus enabling effective low-dimensional mean-field descriptions. In contrast, the amounts and complexity of data produced in the sciences and society is staggering and continuously increasing. In response to this development, since the turn of the century we have witnessed increased research activity on networks, i.e., graphs. They have become fundamental as efficient representations of data. We will therefore carefully discuss networks with and without (non-linear) dynamics. We will do so without necessarily restricting ourselves to small networks since we need tools to disentangle causality in large non-linear networked systems. Here, at the intersection between structure (networks) and dynamics (the equations describing the unfolding of the mechanism behind a graph), is a key frontier where challenges to understanding causality converge. For example, why does this neural network succeed or fail at a classification task, what are the causes that produce the data originating from a networked system? Advancing from observations of data to a causal understanding of what goes on behind the scenes, coming to grips with the generative processes producing observed phenomena, remains a fundamental challenge in science. Paraphrasing Shakespeare, if the world is a stage, then what is the play, or what is the program or code generating the dynamics we observe playing out on the stage? In other words, how

https://doi.org/10.1017/9781108596619.002 Published online by Cambridge University Press

can we find the low-dimensional effective description of a system from observations, i.e., from collecting data? The amount of data makes the stage huge, and generally speaking we do not know the size of an effective low-dimensional description. This is how causality enters the picture. Which state variables are sufficient to capture the dynamics of the system, including different kinds of interventions (causes) and the resulting responses (effects) from this low-dimensional system? In science, we like to ask what effects arise if $X$ occurs, where $X$ could be a natural cause or a human intervention into a system. These are questions we usually ask about the world, i.e., questions about how things work. When dealing with living systems, we like to ask causal questions concerning how such complex systems remain stable and persist (or not) over time. In essence, the existence of regularities in the world, including living systems, implies the existence of manifolds in a geometrical mathematical sense, thus bringing us back full circle to Poincaré's original insights. Here, we'd like to further remark that living systems, such as the human brain, for example, could be viewed as devices evolved to decipher the causal structure around us. Hence causality is not only an ontological question having to do with how the world works; the challenge of causality also has an epistemological component, as it appears to be central to how living systems acquire useful knowledge about the world, facilitating their survival. Thus, what appears to be purposeful behaviour is akin to a teleological aspect of living systems, teleology being Aristotle's attempt to address the issue. Regardless, our focus is to fundamentally ask whether – given observational data – we can reconstruct the generative (causal) program producing a string of events. This is relevant for understanding complex systems, including living systems, where as a rule we do not have prior access to the generative equations. We note that by building powerful systems able to decode such generative models from data, we may gain insight into how living systems acquire and act upon their internal representations of regularities in the world.

In summary, there is a dual sense in which causality is central to understanding living systems: understanding their organisation, which evidently is a composite of complex non-linear systems, and understanding the machinery that enables these systems to extract causal regularities from the world. Now, returning to the topic of stability and persistence as a consequence of the organisation of living systems and the ability to extract causality from the world, it is clear that we need strong fundamental tools to probe causality in such complex systems. Information, in a precise technical sense, provides us with language to assess organisational features of systems and their transformations, without being committed to a particular physical ontology. For example, living (and other) systems can be viewed essentially as information processing entities, thus requiring a language which does not necessarily refer to the underlying physical entities, such as information theory. Classical Shannon information theory is an excellent tool to describe communication channels and transformations, in essence counting bits or symbols given a probability distribution. There are also deep connections with the second law of thermodynamics, for example see [1], as regards to how there is – on average, over time – an increase of entropy, amounting to an increase in the number of micro-states compatible with macroscopic behaviour. There is an immense literature on various applications of Shannon entropy in physics,

https://doi.org/10.1017/9781108596619.002 Published online by Cambridge University Press

biology, and language research. Yet, as we shall see in this book, there are cases where a counting (Shannon) approach is not sufficient to detect differences in structure in systems, i.e., to capture differences in regularities, and therefore by extension, causality. To count bits or states, we need to know what macroscopic property we are interested in. This is a bias, akin to the challenge of optimisation in machine learning; we need to know which function to minimise or maximise. This is not an issue in cases where we know which feature we are searching for. However, it becomes a fundamental restriction when we require a powerful tool to detect regularities in the absence of a prior assumption of an objective function. This, in essence, is what has motivated us to search for a more unbiased quantitative language to describe complex systems and their transformations. Algorithmic information theory (AIT), with roots in precise fundamental theorems in mathematics, is exactly such a framework. The seminal work of Kolmogorov, Chaitin, Solomonoff, Levin and Martin-Löf proved to be sufficiently powerful to enable an unbiased characterisation of randomness and complexity, in effect solving Hume's inference problem definitively. The minor catch is that quantities in AIT are not computable in a Church–Turing sense. This has been a stumbling block with regard to applications of AIT, for obvious reasons. Recent developments spearheaded by our teams and colleagues have demonstrated that these quantities can be numerically approximated beyond the limitations of lossless compression (which we also proved to be almost exclusively related to Shannon entropy rather than to algorithmic complexity), thus providing upper bounds to the complexity of objects.

In this book, we develop a comprehensive account of these concepts from AIT and their temporal, i.e., dynamical, extensions, thus relating them to dynamical systems and causality, using networks as a conceptual bridge in several cases. Our title, *Algorithmic Information Dynamics – a Computational approach to Causality and Living Systems*, refers exactly to these new developments and how they are practically (computationally) and conceptually related, intertwined, and key to probing living systems from the dual perspectives of external and internal causality. We strongly believe that by facilitating an understanding of all these themes and their interconnections, we can hope to make new breakthroughs across several areas of science and, in particular, in the science of living systems.

https://doi.org/10.1017/9781108596619.002 Published online by Cambridge University Press