

Chapter 2

From the Big Bang to Living Cells



We have seen that networks are complex systems, whose complexity, i.e. internal structure, arise from storing information. Living systems are information-storing devices, but so are other similar systems, like ant colonies or human societies.

Because we want to study the aggregation of matter from the Big Bang to the emergence of Artificial Intelligence, this chapter starts with a definition which encompasses all systems we will encounter in the book.

We then go on revising some theories on the emergence of life. Here we will see that molecular biology has identified communication as the engine behind the creation of the first biological cells.

Intelligent Systems

Like a huge fractal, where each part is similar to the whole, life on earth appears to develop by subsequent degrees, creating aggregates of aggregates of said systems. It starts from proteins (aggregates of amino acids) through cells (aggregates of proteins) and complex organisms (aggregates of cells) to social networks – the super-organisms (Hölldobler and Wilson 2009) of ants and, lately, humans.

Internet isn't a living biological network, but it was created by biological organisms, exactly as biological cells created nervous systems. All these systems have the following properties which we observe from the outside:

- they're systems able to **store and process information**. The capacity of managing information is an emergent property – memories are systems in which information is recorded by connecting elements (proteins, cells, people, computer). The whole is more than the sum of the parts: therefore they are also *complex*, according to Simon's definition.

- they are **self-sustaining**: they can nourish their complexity autonomously. They have internal adjustment mechanisms, and manage to absorb sufficient energy from the external environment to maintain their internal structure and grow.
- they are **self-organising**: “A self-organising system is one in which a structure appears without specific intervention from the outside” (Griffith 2013). This concept derives from physics, but as the quote taken from a book on educational theory suggests, it’s now commonly used.

There’s no term in the theory of networks (see Appendix 1) to describe similar systems. We’ll call them **intelligent systems**, with reference to the etymological origin of intelligence, from the Latin *inter-legere*: to join together, connect, in which *legere* has the proto-Indo-European root L-Gh that can also be found in the Italian word *legno* (wood), in the English word *log*, and in the Greek word *λόγος* (logos, word, speech, reasoning). The origin of the concept of *intelligence* is the ability to gather firewood, an essentially human activity used to extract energy from the environment: there are no known *Homo sapiens*’ communities that don’t use fire, and no other species of animals are known to use it. Being intelligent means *knowing* how to extract energy from the environment to remain self-sustaining and self-organising.

“Living networks” could also be a good definition: but not all intelligent systems can be considered biologically alive, although all can be traced back to biological life.

This race for ever-greater complexity, more information processed, more energy consumed, appears to be as relentless as it is necessary. Can we then understand why these systems emerge, and why they seem to have to evolve?

Ex-Nihilo Energy and Information

Before we analyse the evolution of intelligent systems on earth, let’s take a look at the creation of the universe – because all things considered, it’s the structure of the universe that’s responsible for the existence of our solar system and therefore also planet earth. Furthermore, the concept of information we considered in Chap. 1 is of late being seen as a possible key for interpreting the creation of the universe itself.

Every culture has its own creation myth. “Ours” (the scientific one) appears to be by far the most complex. But it has two big advantages.

First, it explains a lot of what we observe – not simply what we see when gazing at the stars on a dark night, but also what we see when we make protons collide close to light speed, and what we measure, with manic precision, with reference to the evolution of the universe.

For example: background radiation. In the universe, radiation appears to come from all directions. If you aim a satellite-TV dish tuned to 160 GHz in any direction, you’ll receive an almost uniform signal. The two physicists who discovered this phenomenon by chance in 1964 won the Nobel prize, as they deserved to, because

years earlier the existence of background radiation had been predicted as part of the Big Bang theory, in exactly the same way in which it was later measured.

There aren't many civilizations that aim satellite dishes at the night sky hoping to receive electromagnetic signals, and there are even fewer creation myths that can explain the existence of this radiation.

The second advantage of our creation myth is that it's based on very simple principles: interactions between matter are explained by so-called spontaneous symmetry breaking. As the temperature of the universe drops, asymmetry appears in the form of new forces.

Spontaneous symmetry breaking in quantum field theory isn't easy to understand, but the phenomenon that's usually taken as an example, is. Imagine a red-hot ball of ferromagnetic material. At high temperatures, the material won't be magnetised. As the temperature drops, the ball becomes a magnet, with the magnetic field in some random direction.

This means that by lowering the temperature we can *acquire information* on a system of balls. In a system consisting of many balls, all above critical magnetisation temperature, we cannot distinguish one ball from another. Below the critical temperature however, each ball will be different to the next: as the direction of the magnetic field of a ball is casual, so each ball can be identified by its magnetic direction. In practice, the difference between macrostate and microstate collapses: from maximum it becomes minimum, and the system "lets us" describe it in a better way, with more information, i.e. less uncertainty.¹

Similar mechanisms lie behind the laws that have governed the universe since its primordial explosion, the Big Bang, 13.8 billion years ago. The most interesting thing is that in the last 50 years, in a more and more decisive way, temperature (in other words energy density) and information have been used to probe the "metaphysical": the very origins of matter-energy and time that constitute our universe (Tryon 1973), (Vilenkin 1982), (Wheeler 1990), (Lincoln 2013).

Information and energy, in these theories, are not linked merely on a thermodynamic level, but at a much more intimate level: energy represents the other side of the coin to information. Having information means having energy, and it's *the spontaneous creation of information that represents the basis for the emergence of matter in the origin of the universe*.

Although the creation of the universe remains within the scope of metaphysics, one of the merits of these speculations is interpreting energy in terms of information, and therefore probability: something that's even more fundamental than energy itself. Probability, in fact, seems independent of the type of universe in which we happen to live: we could imagine a universe in which the fundamental interactions are different, but it's difficult to imagine a universe in which the probability of coming up "heads" when tossing a fair coin isn't 0.5.

¹As we've already seen, entropy should be considered a property of certain macroscopic variables (those that describe the macrostate), and information is always relevant information between two systems (Rovelli 2015)

The Emergence of Complexity

In Chap. 1 we saw how complex systems become all the more complex as the amount of information they contain increases, but the first great complex system actually emerged spontaneously: is the universe, or cosmos, rightfully from the Greek word κόσμος (kosmos, order).

Around 300,000 years after the Big Bang, the universe was an immense, rapidly expanding cloud of hydrogen, with just a little helium and a few other elements. We can compare this gaseous mass to a random network (Erdős and Rényi 1960): each atom influences and is influenced by only a small, almost constant number of other atoms.

In practice, there is maximum entropy in the system, because all the elements – the hydrogen atoms – are indistinguishable inside the network (see Appendix 1 for how to compute the entropy of a network). It seems like the end of the universe: as every system slips towards greater stages of entropy, and as the universe is now at a stage of maximum entropy, it should remain as it is and not evolve.

In reality, a system cannot become ordered unless it is supplied with energy. And although the universe cannot receive external energy, it is a system in unstable equilibrium. The entire universe cannot “recollapse” in a Big Crunch because of the explosive power of the initial Big Bang, but areas in which the gas becomes more compact can form on a local scale.

Imagine the explosion of a bomb made up of magnetic parts: even if the average distance of all the parts from the centre of the explosion continues to increase, some parts will still come together. In practice, instead of many small parts moving away from the centre of the explosion, there will be fewer, bigger ones.

Something similar occurs in the universe. The gas cloud collapses in various points. High-density gas aggregates under extremely high pressure and forms the stars. Stars are indeed compressed hydrogen that fuses into heavier elements in the cores. The fusion produces heat, which stops the stars from collapsing.

The first effect of the formation of stars is the creation of thermal gradients – areas at high temperatures (the stars) and low temperature (everything else). One can extract energy from such a system with any thermodynamic engine, i.e. something which extract heat from the hot reservoir and releases it to the cold one, transforming part of it in mechanical engine.²

The second effect is the **spontaneous emergence of complexity**. Not all stars are created equal: some are very rich in mass, heavy (few) and others are very poor, light (many). In other words the size of stars has the same structure of the wealth of people in a society does, the so-called Pareto distribution (or power-law, see Appendix 1). Furthermore, stars form aggregates, galaxies, in which the number of stars follows the same distribution (few really big, many very small), and galaxies form aggregates of galaxies (clusters) with the same Pareto distribution.

²Steam and internal combustion engines work in this way, see for instance https://en.wikipedia.org/wiki/Heat_engine.

In short, something incredible is happening: matter organises itself and the structures typically found in intelligent systems appear: Pareto distribution and self-similarity.

The emergence of the Pareto distribution in this case has nothing to do with the process of storing information – it’s a “cosmic coincidence” (Watson et al. 2011). When a glass breaks on the ground, the size of the pieces also follow the Pareto distribution, but there’s no form of intelligence behind the process (see Appendix 1 for why this happens).

Intelligent process or not, the universe has now a structure which makes possible the creation of more complex matter, beyond simple hydrogen and helium. In big stars, gas is subject to greater pressure, and burns faster. The few macro-stars burn all their fuel in a few billion years, and explode. Just before the final explosion they produce all the elements in the periodic table. Supernovas, as they are called, have been throwing elements out into space for 12 billion years. One supernova is also at the origin of the matter from which the planets in our solar system are made, and therefore us too.

In physics, the model that describes the formation of all elements – from hydrogen to uranium – during the life of the universe, was developed in the 1950s, and called B²FH, after the authors’ initials (B²FH 1957).

The prodigiousness of the B²FH model is that it not only explains what we observe in the universe. It provides us with an image of a universe in continuous change also in chemical terms. New elements continue to be formed with the passing of time, and the abundance or scarcity of elements lets astrophysicists calculate the age of various regions in space. It is estimated for instance that our solar system, earth included, formed 4.56 billion years ago.

That’s when life could emerge.

Life Without Selection

Charles Darwin stated that asking questions about the origin of life was the same as asking questions about the origin of matter. This did not stop him imagining scenarios which, he thought, might have favoured the formation of living organisms: “some warm little pond with all sort of ammonia and phosphoric salts” (Peretó et al. 2009).

Darwin was even criticised by his admirers for not having proposed an official theory on the origin of life (Peretó et al. 2009). We must give him credit for having avoided taking up the challenge. In the *On the Origin of Species*, Darwin described the evolution of a few complex organisms, made up of biological cells, the structure of which was unknown at the time. It was before Gregor Mendel discovered the laws of inheritance (1865), and of course before DNA was even imagined.

Darwin probably didn’t propose a theory for the origin of life simply because applying Darwin’s mechanism of natural selection to the emergence of life, as done by Dawkins (1976), is like comparing apples with pears (Johnson 2010). What’s more, the idea that a self-replicating molecule with an information content *casually*

appeared in a primordial soup, as imagined by Dawkins (1976) (“At some point a particularly remarkable molecule was formed *by accident*. We will call it the Replicator.”) appears to be statistically groundless (Yockey 1977).³

The fact that the “Replicator” cannot have appeared merely by chance has been considered proof that there must be an intelligent design behind it. As usual, whenever there is no clearly valid scientific explanation, the *intelligent designer* comes in. When a plausible explanation becomes common sense, as in the evolution of the universe, the *intelligent designer* retreats.

Consider a statement like the one John C. Eccles made just in 1989: “Since *materialist solutions fail to account for our experienced uniqueness*, I am constrained to attribute the uniqueness of the Self or Soul to a supernatural spiritual creation.” (Eccles 1989). Today, scientists think that purely materialistic models can provide a perfect explanation for the emergence of conscience (Dehaene 2014), and few people, let alone a scientist, would defend Eccles’ idea.

In a similar way, our knowledge of molecular biology today leads us to consider the creation of life through divine intervention, or the emergence of the same “by accident” as in (Dawkins 1976), not to be the best choice.

The question is therefore: with the instruments at our disposal, can we describe the formation, evolution and behaviour of various intelligent systems, from biological cells up to human societies, going through the nervous systems?

The Startups of Life

There is often some confusion about physics dictating that “it’s impossible to create order” or “disorder is constantly increasing.” We have seen the universe itself created order soon after the Big Bang. When we freeze water in a freezer we create order. It’s just that we need energy to do so.

Evolving systems can be divided into 4 categories:

1. The system becomes more ordered and absorbs energy – the “freezer” system. Possible only if we provide sufficient energy.
2. The system becomes more disordered and emits energy – the “explosion” system. Possible only if there is not too much energy emitted.
3. The system absorbs energy and becomes disordered – the “adolescent” system. These reactions are always possible. It’s easy to waste energy for creating disorder!
4. The system becomes more ordered and emits energy – the “genie in a bottle” system. Unfortunately, it’s impossible.

³There is still the possibility that life appeared on earth after this casual process had occurred an infinite number of times: both in an infinite number of universes and in this universe, and billions and billions of times in systems similar to our solar system. It’s possible, but perhaps it would be better to come up with a mechanism that makes us less unique...

Figure 2.1 shows how the systems above evolve. They all start from a state where no energy is absorbed or emitted, and the order is not perturbed. They then start evolving, and finally stop. The fridge is switched on, absorbs increasingly more energy, creating more order, then stays in a region where it absorbs an almost fixed amount of energy, creating a fixed amount of order. Until it's switched off.⁴

But if we take life since the very beginning, as shown in Fig. 2.2, there are two things which do not fit.

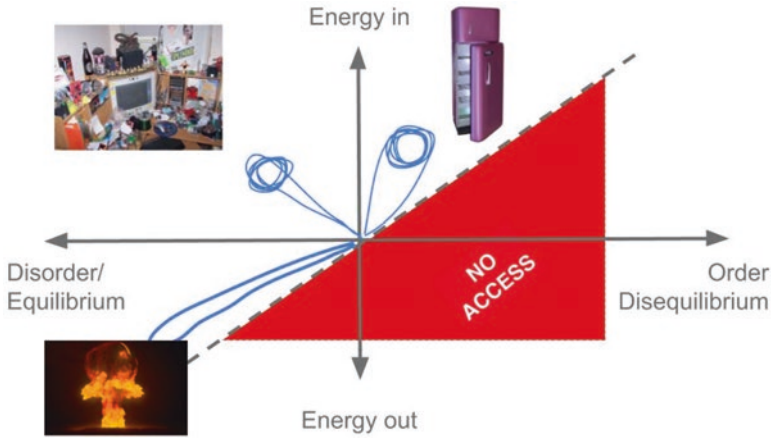


Fig. 2.1 The three possible evolution of a system in the space “energy absorbed/emitted” versus “order created/destroyed”

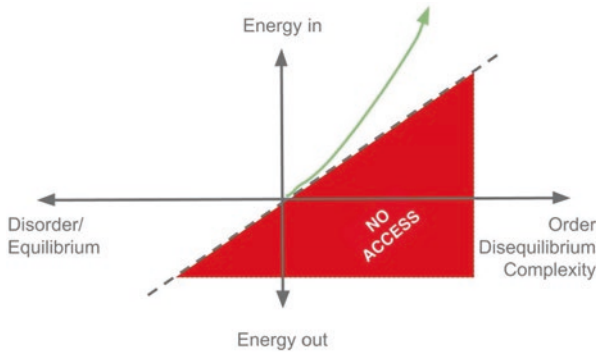


Fig. 2.2 Since its appearance, life has been absorbing increasingly more energy from the environment, constantly escaping equilibrium, i.e. creating more order

⁴Pictures available under creative common licence: <https://www.needpix.com/photo/download/986345/refrigerator-freezer-fridge-freezer-retro-seventies-american-style-metalic-cool-cold>, <https://www.flickr.com/photos/puyo/253932597>, <https://www.flickr.com/photos/51686021@N07/33230171512>

The first is –why it started? A fridge starts cooling, when we provide energy, because that’s the way we built it. We are the intelligent designer.

The second is –why it keeps growing? “Because life evolves” is not a good answer. “Because of natural selection” isn’t either: why should natural selection bring an increase in energy absorbed and order?

To answer, let’s have a look at the ecosystem of companies. Specifically, startups growing without external funding (the ones who can “bootstrap”,⁵ a concept we will encounter again later). A startup which bootstraps has little or no internal organization, just a few elements, but is able to absorb some energy (cash from sales).

If the business model, i.e. the information retained by the founders on how to extract energy from the market, is valid, the startup will use some of the earned cash to increase in size. Together with that, internal organisation –complexity– will appear.⁶ The organisation represents an asset for the company –is what allows the exploitation of the market– but also a liability –it requires energy in the form of cash.

The startup needs little cash at the beginning while providing, with its “disruptive idea”, a way to extract revenues from the environment. The extra cash will be used to improve the internal organisation, while increasing revenues. It’s an avalanche effect: more organisation increases the revenues, bringing more organisation, which requires more cash but allows more revenues and so on.

With growth, there is a scale effect: bigger networks –companies– can process more information per element, notwithstanding the fact that the information processed by each element is less than it was at the beginning. It’s Herbert Simon definition of complexity –the whole is bigger than the sum of its parts.

In a similar way, we must find a mechanism which allows the aggregation of basic chemical elements into something more complex through absorption of energy. The information on how to extract energy must be contained somewhere, and these elements must be able to “read” it. Once the elements get together, they must be able to store information on how to aggregate further, storing increasingly more information.

It looks like an impossible task but, seen that we are here, alive, it must be not.

Amino Acids – The Entrepreneurs of Life

Amino acids are big molecules which, bonded together, form proteins, which eventually organise themselves into biological cells. One of the most interesting discoveries in biology during the last century was that amino acids, considered the building blocks of life, appear spontaneously in nature.

⁵“Building a company from the ground up with nothing but personal savings and, with luck, the cash coming in from the first sales”. <https://www.investopedia.com/terms/b/bootstrap.asp> verified 26 August 2019.

⁶Lex Donaldson, a sociologist, defines complexity as “a measure of the amount of knowledge available within the organization (Donaldson 2001)

Soon after World War II, Harold C. Urey, a Nobel prize and key figure in the development of the nuclear bomb, began to consider the emergence of life in an atmosphere rich in methane, hydrogen, and ammonia (Bada and Lazcano 2003). It's Darwin's "primordial soup".

Urey's idea had such an impact on the 22 year old Stanley Miller, another physicist, student of Edward Teller – often referred to as the Father of the hydrogen bomb⁷ – that he decided to abandon nuclear physics for setting up an experiment which proved Urey's hypothesis.

In 1952, for the first time, humanity produced synthetic amino acids in the Miller-Urey experiment.

Today, we believe primordial conditions were different to those imagined by Urey and Miller, but the fact remains that the formation of amino acids can be a spontaneous process. It's a natural "freezer system", a little order is created and in exchange a little energy is absorbed.

But how amino acids, these big molecules, can aggregate in more complex structures, like proteins? This cannot be a spontaneous process. To bond amino acids we need much more than just electrical discharges: they require specific chemical processes. And these chemical processes do not occur by chance, especially on the necessary scale. In addition to that, even if we managed to supply energy in the right way, we couldn't create long amino acid chains – proteins – as we wished, because they would break.

One mechanism that partially solves the first problem was proposed soon after Miller's experiment (Koshland 1958). There are substances, called enzymes (from the Greek word *ἔνζυμον*, in yeast) which, due to their particular form, and therefore the electric field around them, can favour the creation of a bond between two amino acids, the peptide bond.

It's like building a chain of lockers with or without their keys – energy is required in both cases, but much less in the first case. In presence of enzymes, amino acids bond easily, using just some excess of the energy available in the environment.

There is still the second problem: it's impossible to have proteins, chains of billions of amino acids, free in the environment. Imagine a human chromosome – a DNA molecule – lying unfolded: it would be two metres long. In any natural environment it would break immediately.

This problem is solved by collaborative molecules, called "chaperones". Molecular chaperones can favour, without having to do any work, the folding of the proteins into stable configurations. These stable configurations have the additional property of favouring proteins interaction. When folded, proteins can move in water, interacting with each other without getting trapped.

Egg white is an excellent example of how important protein folding is. When raw, it's fluid, and consists of an aqueous solution of folded proteins. When beaten into a thick foam or boiled, the white solidifies because the proteins burst and are

⁷For his obsession in building such a destructive device, together with his strong Hungarian accent, Edward Teller is also the person the character Dr. Strangelove in Stanley Kubrick's film was based on.

trapped in each other. In the first case the proteins can move around and interact, in the second they are immobilised.

In conclusion, for proteins to form, survive and eventually interact we must have:

1. A mechanism which creates enzyme for amino acids to bond together in proteins
2. A mechanism which creates chaperones, for the proteins to fold

The Secret of Life

Ever since it was discovered, DNA has been called “the molecule of life”. But really, it would be more fitting to call it “the brain of cellular life”. Our brain, while essential for our survival, could not survive alone. The same is true for a DNA molecule, which alone would remain helpless and in time disappear.

Nevertheless, the idea that DNA is key not only for cellular life, but also for the appearance of life itself, has never been far from the thoughts of many chemists and biologists alike: one of the most popular theories on the origin of life is in fact the so-called “RNA world”, in which ribonucleic acid, a molecule similar to DNA, was responsible for having kick-started life.

The “RNA world” hypothesis derives from the fact that, in a similar way to Miller’s experiments, we’ve managed to reproduce the conditions in which amino acids organise themselves and bond in RNA chains which, in the right conditions, start reproducing. The only problem with this theory is that in any case, at the start of the process, you have to introduce enzymes of biological origin. Moreover, DNA and RNA are used by cells to store and transport information on how to construct proteins, so we don’t know how an RNA that reproduces autonomously could evolve into cells (Robertson 2010), (Davies 2000).⁸

Another hypothesis for the origin of life, better-suited to the framework of intelligent systems, is the “proteins first”, proposed by physicist Freeman Dyson (2004). Going against the idea that life started by creating both the information medium and the information itself (RNA) *at the same time*, Dyson took inspiration from the theory of Russian biologist Alexander Oparin (1957). Oparin had the intuition that life is made possible by how molecules interact. Life begins with the interaction between proteins, and not with their capacity to reproduce.

What’s fascinating about Oparin’s view, is that life is seen as a natural process in the evolution of matter. Previously we considered how the universe created more and more complex structures, starting from a shapeless *blob* of energy, the Big Bang. In Oparin’s opinion, life should be considered the continuation of this process – a continuation that stopped at very low levels in some parts of the universe, while on earth and perhaps on other planets it continued to create complex organisms.

⁸“Is Eigen’s work a reconstruction of the steps in which nature created life from inanimate matter? Evidently not. “Page 139 in the Italian version.

The idea that life started thanks to the fact that amino acids spontaneously became organised into complex structures has become all the more credible in recent years thanks to theoretical and experimental works. Guseva et al. (2017), for example, introduce a complexification mechanism based on the polarity of water molecules and polymers.⁹

The mechanism is ingenious, and surprisingly simple: the authors prove that in certain conditions a small percentage of proteins¹⁰ fold in water, because some amino acids are hydrophilic and others are hydrophobic: the first tend to form the surface of the folded protein, while the second form the internal part. Some hydrophobic amino acids remain exposed though, attracting the hydrophobic elements of other proteins.

Sometimes two hydrophobic tails fall together into the field of another hydrophobic section on the surface of the refolded protein. Like a magnet that attracts other magnets so they stick together, this section helps the two tails interact and form a bigger protein – it's the enzymatic process. In practice, small proteins become the enzymes that help other proteins grow ever bigger. The same mechanism helps these new proteins fold.

Some proteins, highly interacting with each other, combine spontaneously in closed communities, called protocells, that can attract or search for the substances they need to function. The fact that a few molecules combined, *protocells*, can form systems able not only to use the chemical energy available in their environment, but also to move in order to find it, has been proven in various experiments (e.g. see Hanczyc et al. (2003), made famous by Hanczyc's TED talk¹¹).

Protocells are dissipative structures: structures that can use an energy flow to create and maintain an internal order (Kondepudi and Prigogine 1998), (Margulis and Sagan 1997). When there is no energy, the internal structure disappears, exactly as an organism without food decomposes.

While remaining within the field of speculation – there is still no widely accepted paradigm for the origin of life – according to Dyson at this point the RNA can form (as there are enzymes) and reproduce (the RNA can produce its own reproductive enzyme, the ribozyme).

The RNA initially appeared as a parasite. After having exploited the spontaneous emergence of simple polymers/enzymes, it started to grow, absorbing the same polymers it no longer needed thanks to the ribozyme, destroying the host protocell.

At this point, the RNA had two options: continue to emerge spontaneously, grow, and destroy every protocell it manages to infect, and therefore disappear, or collaborate with the protocell on which it is the resident. Since 1982, in fact, we know that

⁹A term used generically for chains of simpler molecules, the monomers. It includes also, but not only, proteins.

¹⁰Specifically, we are talking about very small proteins, so small that they are called peptic chains and not proteins.

¹¹Hanczyc, M. M. (2011) The line between life and not-life. TEDSalon London Spring 2011. https://www.ted.com/talks/martin_hanczyc_the_line_between_life_and_not_life

RNA acts in cells not only as an information medium, but also as a protein enzyme: in other words it helps other proteins form.

In the words of Freeman Dyson:

Within the cell, with some help from pre-existing enzymes, the nucleotides produced an RNA molecule, which then continued to replicate itself. In this way RNA first appeared as a parasitic disease within the cell. The first cells in which the RNA disease occurred probably became sick and died. But then, according to the Margulis scheme [which we'll discuss later], some of the infected cells learned how to survive the infection. The protein-based life learned to tolerate the RNA-based life. The parasite became a symbiont. And then, very slowly over millions of years, the protein-based life learned to make use of the capacity for exact replication that the chemical structure of RNA provided. The primal symbiosis of protein-based life and parasitic RNA grew gradually into a harmonious unity, the modern genetic apparatus. (Dyson 2004)

Dyson adds that what his is “not yet a serious scientific theory”. One might add, nonetheless, that the idea that life followed what happens at a macroscopic level – first the medium for storing the information (hardware) is supplied, and then the information itself (software) – appears to be more plausible than the RNA world theory.

The idea that a parasite can turn itself into an information processing centre is incredibly appealing and, far from being a phenomenon specific to cells evolution, it's something we come across in nature –and in our society, as we will see– all the time.

A parasite is an organised structure that grows in another organised structure, in a way that's detrimental to the same. The RNA parasite found its way into cells – fragile structures always looking for new mechanisms to supply the necessary energy – finally destroying both itself and the cell.

It would be much better, for the protocells, to learn how to use this fantastic molecule, and use it on the one hand as an enzyme, and on the other as an information medium to memorise which proteins were necessary, and when.

With the inclusion of RNA also came the ability to replicate, and therefore also evolution through natural selection. Selection, for the reasons explained above, cannot lay the foundations of life; it's just one of the processes that allows life to evolve. As the physicist Jeremy England said, “you might call Darwinian evolution a special case of a more general phenomenon” (Wolchover 2014).

The secret of life is collaboration, not selection.

Evolution Through Learning

Even if we accept that proteins managed to absorb an RNA molecule and use it as a data processing centre, one fundamental question remains unanswered: how did protocells with RNA evolve into cells able to duplicate their “brain” and pass it down to descendants?

So we must imagine a mechanism that drives life towards a stage of *evolvability*, or a propensity for evolution (Watson and Szathmáry 2016).¹² For example, the

¹²See “Outstanding Question 1”

ability to reproduce is essential for the evolution of organisms. By reproducing, an organism increases its *evolvability* also thanks to natural selection, and its code will be favoured over others (reproduction also favours evolution because, as we'll see below, makes it also possible to reset unnecessary complexity).

We saw how Hebbian learning works in the brain (Chap. 1, *Storing information*): two neurons that are activated together reinforce their connection. The same thing is true for language, and it's a classic example of unsupervised learning in Artificial Intelligence: if the words "kitten" and "cat" are often used in the same context, a computer (or a person) understands they are synonymous – i.e. they are connected – without needing external training.

But we've also seen that this mechanism is the one many complex networks use for memorisation. Not only neural networks, but also societies (two people who have the same interests are likely to form a bond) store information as Donald Hebb described in the 1940s. Creating paths between the elements, and assigning a meaning to each path.

The same thing probably happens with genes. Exactly like the brain doesn't memorise in terms of single neurons (the "words"), but rather in terms of the connections between the neurons (the "phrases"), a cell's information is stored in a genetic network, the GRN (Genetic Regulatory Network) (Davidson et al. 2002)

The cell learning mechanism is essential for adaptation in unicellular organisms, and is used for the formation of complex organisms. How can two cells with exactly the same DNA become part of such different parts of the body as the intestine and the brain? By "learning" what they must become. Learning doesn't necessarily have to be Hebbian, as there are non-Hebbian mechanisms that let the cell inhibit or favour certain genes (Epigenetics 2013), but Hebbian learning was recently shown to be a possible learning mechanism (Watson et al. 2010), (Watson and Szathmáry 2016).

Artificial Neural Networks and DNA

The Hebbian model was also the model that inspired the first attempts to replicate the learning capacities of biological neural networks with computers, although without much success (Anderson and Hinton 1981).

The situation changed in 1982, when physicist John Hopfield proved that it was possible to define an "energy" function for mathematical structures resembling a biological neural network.

Physical systems tend to settle into low energy states, like a ball in a bowl. If we have a function describing the energy of a system, we can also identify stable configurations. This is what Hopfield did: he proposed a function computing, for any configuration, a certain variable, which he called "energy". Starting from an input, it was possible to find, always, a particular configuration minimising this "energy" function. Similar inputs would have the same minimum, and therefore would correspond to the same "stored message". It's like saying we can recognise similar bowls because the ball will move to the same position however we throw it in (Hopfield 1982).

In practice, Hopfield networks are defined by the weights of the relationships between the neurons. The energy of a neuron is the sum of the weights for its state of excitation (0 or 1). The energy of the network is the sum of the energies of each neuron. When each neuron tries to minimise its energy value, independently of the others, a Hopfield network falls into a configuration (neurons activated / deactivated) that minimises the energy of the whole network.

To store a vector [a list of 0 and 1s], the weights are changed to reduce the energy of that vector. So the stored vector correspond to the local minima in the ‘energy landscape’. To retrieve a vector from a noisy version of it, the network is put into an initial state that represents the noisy version and then allowed to settle to an energy minimum (Hinton 2014).

But what have Hopfield networks – a computational neuroscience instrument – got to do with the genetic code? A lot! The genetic code shouldn’t be seen as a mere *storage device*, dumb memory. It’s more similar to our brain than to a hard disk. Paul Werbos, the mathematician who in his 1974 graduate thesis showed how we can “make artificial neural networks learn”, wrote:

...to what extent is the “junk DNA,” 97% of the genome, actually a kind of learning system, like the brain. More precisely, how much of the genome is intended to help us learn how to choose better gene expressions, as opposed to merely specifying the final “actions” at the output layer of the “gene brain”?¹³

Geoffrey Hinton, probably the most important figure in the development of artificial neural networks, as long ago as 1987 wrote, together with Steven Nowlan:

Many organisms learn useful adaptations during their lifetime ... It seems very wasteful not to make use of the exploration performed by the phenotype to facilitate the evolutionary search for good genotypes. The obvious way ... is to transfer information about the acquired characteristics back to the genotype (Hinton 1987).

While today epigenetics have legitimised the idea that genetic code can pass information acquired during the life of the organism on from parent to offspring, this wasn’t so obvious in the 80s, and John Maynard Smith – one of the most important geneticists and scholars of evolution of the century – had to publicly support Hinton’s (1987) publication when this was blocked from publishing (Smith 1987). The concept, taken up again by Watson and Szathmáry (2016), is clear: to evolve, a successful system must be able to learn, and pass what it learns on to its offspring.

A living system which, when it dies, loses all the information it acquired during its lifetime would not only pointlessly destroy something precious, but also slow the evolution of the species to such an extent that it would soon become extinct (Hinton 1987).

According to biologist Lynn Margulis (1997), to whom next section is dedicated, a crucial feature of living systems is that the internal organisation they support is simply the information necessary for finding the energy they need – exactly what we discussed in the previous chapter.

¹³<https://www.werbos.com/life.htm>, retrieved August 10, 2018.

In a more general framework the ultimate purpose of every intelligent system, including living systems, is to reduce uncertainty concerning the external environment, storing information, to acquire the energy necessary to sustain its internal order. It's therefore only natural that life has adopted a data processing centre, DNA, that lets it learn, and therefore quickly evolve.

Collaboration and Eukaryotes

The Dyson quotation in the previous section refers to the “Margulis scheme”. Dyson explains that he got the idea from a theory created to explain the emergence of more complex biological cells, called eukaryotes, published by Lynn Margulis (1970).

Dyson (2004) also wrote a short, excellent description of Margulis's original theory:

...the evolutionary tree has three main branches representing a divergence of cell types far more ancient than the later division of creatures into animals and plants. Moreover, the genetic apparatus carried by organelles such as chloroplasts and mitochondria within eucaryotic cells does not belong to the same main branch of the tree as the genetic apparatus in the nuclei of eukaryotic cells. The difference in genetic apparatus between organelles and nucleus is the strongest evidence confirming Lynn Margulis's theory that the organelles of the modern eucaryotic cell were originally independent free-living cells and only later became parasites of the eucaryotic host.

According to this theory, the evolutionary success of the eucaryotic cell was due to its policy of free immigration. Like the United States of America in the nineteenth century, the eucaryotic cell gave shelter to the poor and homeless, and exploited their talents for its own purposes. Needless to say, both in the United States and in the eucaryotic cell, once the old immigrants are comfortably settled and their place in society is established, they do their best to shut the door to any prospective new immigrants.

With his reference to politics, Dyson suggested that the Margulis scheme – the idea that parasitism and symbiosis are the driving forces behind evolution – could help understand not only the emergence of life but also properties of human societies which should be analysed in terms of emergence properties.

The Importance of Scientific Revolutions

When Lynn Margulis died in 2011. The obituary published by Nature praised “her paradigm-changing book, *Origin of Eukaryotic Cells*”, published in 1970 (Lake 2011).

The term *paradigm* refers to Thomas Kuhn's (1962) *The structure of scientific revolutions*. In his work, Kuhn, a physicist and historian of science, proposes a social vision of the scientific method. There are no strict rules on what is scientific and what is not (contrary to what Karl Popper (1934) thought). It's a *shared idea*,

which Kuhn calls *normal science*, that guides a scientist's work. There are works, in this normal science, that can be considered paradigms (from παρα-δείκνυμι, show to the side). Aristotle's *Physics* in this sense is no less a scientific work than Galileo's *The Assayer*, Newton's *Principia*, or Freud's *The Interpretation of Dreams*. All of these works coagulated the knowledge of the time in order to refer to a multitude of concepts, connected with each other, without having to explain what the author is referring to all the time. The concept of force in physics was crystallised by Newton, that of the unconscious by Freud. Every time we use these words, we're referring to the Newtonian and Freudian *paradigm*.

As Feyerabend (1989) then noted, a shared paradigm can eventually stiff the evolution of scientific thought. There will be scientists who'll defend it even if it's indefensible: Aristotelians, like the character of Simplicio in Galileo's *Dialogue*, are the problem, not Aristotle.

The anatomist who in the *Dialogue* observes that all nerves originate in the brain but remains convinced they originate in the heart because Aristotle said so, strikes a sour note ("You've shown me this thing which is so open and sensible, that if it weren't for Aristotle's text of the opposite opinion, as it openly says, the nerves originate in the heart, one would have to confess it was true").

But considering what we have discussed about uncertainty and coins in Chap. 1, the anatomist behaviour might be justified. As we've seen, considering a coin as fair, i.e. that it will have a 50% probability¹⁴ of coming up heads or tails, is the safest choice we can make. We have held thousands of coins in our hands, and none of them seemed biased. If, without an absolutely valid reason, you were to consider the coin biased with an 80% probability of coming up heads and a 20% probability of coming up tails, you'd risk increasing your uncertainty and if taking bets you'd lose money fast: the bet placers would continue to consider the coin fair, and rightly so, while you would take ridiculous bets (such as "heads" 3-1) and lose time and time again.

On the other hand, if the coin is actually biased but everyone, including you, thinks it's fair, no one would win much as everyone has the same average uncertainty. Likewise, if the consensus for 2000 years was that nerves originate in the heart, stating that the opposite is true means risking everything. Just before *Dialogue* was published, Jacopo Berengario of Carpi, the first ever neurosurgeon, was called to operate on the fractured skull of Lorenzo de' Medici, one of the most powerful men in Renaissance Italy. If Berengario had used Aristotelian medicine and Lorenzo had died, his reputation would have suffered. If Lorenzo had died while being operated on by a non-Aristotelian surgeon, as Berengario was, he would have lost not only his reputation but also his head.

The problem was that if he had used the Aristotelian nerve model, the whole society would have lost. The "anti-information" on the origin of nerves caused the deaths of many patients, and it was down to scientists like Galileo and the likes to

¹⁴We always refer to the Bayesian definition of probability: "Probability ... serves to express, in a precise fashion, for each individual, their choice in their given state of ignorance" (de Finetti 1970)

put forward counterarguments convincing enough to destroy the old paradigm and create a new one.

Society will benefit from a new paradigm that minimises uncertainty, and can create new instruments to sustain its internal complexity. It's the same mechanism Margulis (1997) described for the evolution of cells, and can be applied to her own work.

Tenaciously sticking to the paradigm can be counterproductive. For example, using the principle of natural selection as the driving force of evolution of societies – as Darwin himself and many others did – leads to make hazardous conclusions. According to the English naturalist, hospitals and health services were the cause of the decline of our species: “We ... do our utmost to check the process of elimination; we build asylums for the imbecile, the maimed, and the sick; we institute poor-laws ... Thus the weak members of civilised societies propagate their kind ... *this must be highly injurious to the race of man*. It is surprising how soon a want of care, or care wrongly directed, leads to the degeneration of a domestic race; but excepting in the case of man himself, hardly any one is so ignorant as to allow *his worst animals to breed*.” (Darwin 1874).

Darwin was wrong. There was a period, after the end of the Second World War, in which the human race and its living conditions improved, not thanks to selection but instead thanks to social policies and collaboration (Pinker 2011), (Norberg 2016). Fascists and Nazis, who followed social darwinism and wanted to suppress what they considered the “weaker” elements of society, produced as a result hellish conditions for human beings.¹⁵

Margulis proved that life emerges as a result of collaborative processes: collaboration is the driving force behind evolution.¹⁶ This is what the present book wants to take to extremes: collaboration – in practice the exchange of information/energy – lies behind every successful long-term enterprise, from eukaryotes to nation-states. And that's not all: the union of various elements is what makes it possible for everything to store information. If humanity was composed of only the “fittest” it would fail, as would a language made up of only particularly eloquent words. Obviously, every new word can only enrich the language.

What's more, as mentioned in (Margulis 2008), Darwin's *On the Origin of Species* isn't about the origin of life or even the origin of the species.¹⁷ Not acknowledging that natural selection plays a role in the origin of life doesn't make it any less important for the evolution of the species, but “pure” Darwinists seem unable to accept other mechanisms in biology, sociology and whenever they believe selection

¹⁵With the help of a Darwinian foreign policy aimed at defending living space, *lebensraum* (Ratzel 1901), rather than collaborating with other states.

¹⁶The fact that Margulis was worthy of an obituary in *Nature*, which called her book paradigmatic, therefore measuring it with the same yardstick as Darwin's work, did not stop her book from failing to be acknowledged by anyone outside a close circle of scientists, and, unfortunately, being literally eradicated from human memory. No further editions of the book were ever printed and finding an electronic version is impossible.

¹⁷As the subtitle of Darwin explains: it's about “the Preservation of Favoured Races”

could be applied. Let's not forget that Darwin broke ties with the past by valiantly defending the concept of evolution: it's not natural selection that's changed the way we think about life on earth, but rather the idea that life evolved.

If this change occurred thanks also to collaboration and not just selection, this means that a collaborative societies, whatever Darwin thought, might not only be nicer to live in, but also more likely to be successful.

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

