



# How Life Works

A USER'S GUIDE TO  
THE NEW BIOLOGY



**Philip Ball**

# How Life Works

**A User's Guide to The New Biology**

**Philip Ball**

The University of Chicago Press  
Chicago

*[OceanofPDF.com](http://OceanofPDF.com)*

# Contents

Prologue

**1** The End of the Machine: A New View of Life

**2** Genes: What DNA Really Does

**3** RNA and Transcription: Reading the Message

**4** Proteins: Structure and Unstructure

**5** Networks: The Webs That Make Us

**6** Cells: Decisions, Decisions

**7** Tissues: How to Build, When to Stop

**8** Bodies: Uncovering the Pattern

**9** Agency: How Life Gets Goals and Purposes

**10** Troubleshooting: Rethinking Medicine

**11** Making and Hacking: Redesigning Life

*Epilogue*

*Acknowledgments*

*Source Notes*

*Bibliography*

*Index*

## Prologue

On June 26, 2000, US President Bill Clinton announced that scientists had completed a first draft of the human genome. That's to say, they had deduced the sequence in which nearly all of the three billion chemical building blocks of our DNA are strung together. "Today," he said, "we are learning the language in which God created life."

He was wrong, but not (just) in the way you might think.


People are, of course, used to politicians saying wrong things (and not just about science). Yet the two scientists at hand did not rush to correct Clinton. On the contrary, one of them—Francis Collins, then head of the US National Institutes of Health, and now science adviser to President Joe Biden—went on to echo the same sentiment by celebrating this newfound ability to read "our own instruction book, previously known only to God."

Many scientists will have bristled at these religious references, but truly that was not where the problem lay. (At least, not unless you are an atheist or a theologian.) The metaphors of the "language of life" and the "instruction book" of humankind are even today routinely used to refer to the human genome, which was analyzed (almost) in its totality by the international Human Genome Project (HGP) as well as by the privately funded parallel effort run by biotechnological entrepreneur Craig Venter, who was also present at the unveiling ceremony and is an avowed atheist.<sup>1</sup>

More than two decades later, the information supplied by the HGP consortium, and by the subsequent sequencing of tens of thousands of individual human genomes, is proving to be a vital resource for biomedical research. That was always the hope, and a significant part of the mission. But not only has this information brought us little closer to understanding life itself; it has in some ways shown us that we are further away from such understanding than we thought. For if there is anything like a language of life, it will not be found in the genome—which does not resemble any instruction booklet ever made by humans.

Yet misleading metaphors for the genome remain as persistent and popular as ever. The “blueprint” is a favorite, implying that there is a plan of the human body within this three-billion-character string of “code,” if only we knew how to parse it. Indeed, the whole notion of a “code” suggests that the genome is akin to a computer program, a kind of cryptic algorithm that life enacts. The “book of life” has even been given a physical realization: it comprises a total of 109 distinct books, collected into 23 volumes (one for each of our chromosomes), in which page after densely spaced page are filled with the sequence of four letters (a, t, c, g) that represent the building blocks of DNA (fig. 0.1). I am happy to leave the reader to judge which book—that one or this one—offers a clearer picture of how life works. The aim of this book is to show why these metaphors are inadequate, why they need replacing, and why we will not understand how life works until we do. It also attempts to sketch out what might be put in their place.

There’s no shortage of alternative metaphors for the genome itself: it has been likened to a musical score, for example, or the script of a play. Some of these analogies are improvements, though none is perfect. But the key point is that looking to the genome for an account of how life works is rather like (this simile is imperfect too) looking to a dictionary to understand how literature works.



**Fig. 0.1** The “book of life”? The human genome as recorded in 109 books produced from the Human Genome Project. Books made by Kerr/Noble. Image courtesy of the Wellcome Collection.

When biologists are challenged about why the decoding of genomes—ours and those of many other species—has offered so little real insight into the process we call life, they will typically say that it has all proved to be rather more complicated than we had anticipated. As Dutch biologist Bé Wieringa said on his retirement in 2018 after a career devoted to studying how genes affect life and health, “[after the HGP] we thought we’d be done. The reality, of course, is we’re not. In fact, the possibilities have expanded even further.”

Wieringa added rather poignantly, “If I’m honest, I really did believe that cells and molecules [like genes and the molecules they encode] had a slightly simpler relationship.” We all did; the HGP was largely predicated on that belief. Ironically, the project itself has turned out to have offered one of the best reasons why we should relinquish such dreams of simplicity.

But the alternative is not necessarily to capitulate to the bewildering confusion to which Wieringa seems resigned. Instead, the findings of the HGP are an invitation to say “Of course it is not that simple! How could we have ever imagined that life itself could be? But what glorious, subtle, *useful* ingenuity we are finding in its place!”

Letting go is hard, however. The “instruction book” view of the genome persists precisely because the real story about how DNA and other molecules produce and sustain cells and organisms is not that simple. The metaphor offers consolation: it suggests a tidy tale that, even if it is wrong, seems preferable to muttering “Actually, it’s more complicated than that.” And it’s true that once you relinquish the idea that the “secret of life” lies in the genome—if only we knew how to interpret it—biology can look totally baffling. As I will show, just about all the neat stories that researchers routinely tell about how living cells work are incomplete, flawed, or just totally mistaken.

All the same, I believe we *can* do better. I will show how research in molecular and cell biology over the past several years has painted a richer and much more astonishing picture than that bleak and obsolete mechanical metaphor. The picture does at times appear fantastically baroque and perplexing, but in the end it takes the burden of control off the shoulders of the genome, relying instead on principles and processes of self-organization that, precisely because they have no need of tight genetic guidance, avoid the fragility that would engender. I must stress that there is nothing in this new view that conflicts with the neo-Darwinian idea that evolution shapes us and all other organisms and that it depends on the genetic transmission of information between parent and offspring. However, in this new view genes are not selfish and authoritarian dictators. They don’t possess any real agency at all, for they can accomplish nothing alone and lack a capacity for making decisions. They are servants, not masters.

Fundamentally, this new view of biology—which is by no means complete, and indeed is still only nascent—depends on a kind of *trust*. You could say that genes are able to trust that there are processes beyond their capacity to directly control that will nonetheless allow organisms to grow and thrive and evolve. (Biologists need to develop that trust too.) This way of working appears repeatedly in biology when things get complicated and tasks get hard. When organisms first became multicellular, when they

became able to adjust to and exploit the full richness of their surroundings through sensory modalities like vision and smell, when their sensitivity and receptivity to the environment became genuine cognition, it seems that life increasingly relinquished a strategy of prescribing the response of the organism to every stimulus, and instead supplied the basic ingredients for systems that could devise and improvise solutions to living that are emergent, versatile, adaptive, and robust.

The new picture dispels the long-standing idea that living systems must be regarded as *machines*. There never has been a machine made by humankind that works as cells do. This is not to deny that living things are ultimately made of insensate and indeed inanimate molecules: we need no recourse to the old idea of vitalism, which posited that some fundamental and mysterious force made the difference between living and inert matter. Yet dispensing with the machine view of life allows us to see what it really is that distinguishes it from the inanimate world. The distinction is as fundamental and wondrous as the formation of the universe itself—but more amenable to scientific study, and for that reason probably more tractable.

In particular, life is not to be equated with that special kind of machine, the computer. It is certainly true that life performs kinds of *computation*, and indeed there are key features of biology that can be fairly well understood using the theory of information developed to describe modern information technologies. What is more, a comparison with machines can sometimes be a useful way of thinking about how *parts* of the process that is life operates. I will occasionally make such parallels. It is meaningful to say that our cells possess pumps, motors, sensors, storage, and readout devices. That, however, is very different from the modern trend of discussing the fundamental features of living organisms by comparing them to electrical circuits, computers, or factories. No computer today works as cells do, and it is far from clear that they ever will (or that this would be a good way to make a computer anyway). There is so far no technological artifact that provides a good analogy for living systems. These are a different kind of entity, with their own logic, and they have to be their own metaphor.

We are already somewhat familiar with this logic. We know that, to solve difficult challenges, it is sometimes best not to seek a particular,

prescriptive answer by reductive means, but instead to give people relevant skills and then trust them to find their way to an effective solution—one that can be altered and adapted as circumstances dictate. We can now see that by organizing our human systems this way, we are simply reenacting at another level of the biological hierarchy the process already operating within us: we are utilizing the wisdom of how life works.

Central to this new view of life is a shift in the notion of what life itself is. The problem of defining “life” has bedeviled biology throughout its history, and still there is no agreed resolution. But one of the best ways to characterize living entities is not through any of the features or properties usually considered to define it, such as replication, metabolism, or evolution. Rather, living entities are *generators of meaning*. They mine their environment (including their own bodies) for things that have meaning for them: moisture, nutrients, warmth. It is not sentimental but simply following the same logic to say that, for we human organisms, another of those meaningful things is love.

One key reason for the failure of the machine analogy is that cells work at the scale of molecules, and things are different in the molecular world. They are noisy, random, unpredictable—and life does not so much battle to maintain order in the face of those influences as find ways to put them to good use. Life thrives on noise and diversity, on chance accidents and fluctuations. It simply couldn't work otherwise.

There is, then, no unique place to look for the answer to how life works. Life is a hierarchical process, and each level has its own rules and principles: there are those that apply to genes, and to proteins, to cells and tissues and body modules such as the immune system and the nervous system. All are essential; none can claim primacy. As Nobel laureate biologist François Jacob wrote, “There is not one single organization of the living, but a series of organizations fitted into one another like nests of boxes or Russian dolls. Within each, another is hidden.”

Thus, as Michel Morange, a professor of biology at the Ecole Normale Supérieure in Paris, has said, “Biological function emerges from the complex organization that spans the whole scale of life, from molecules up to whole organisms or even groups of organisms. Complex functions find their origin and explanation in this hierarchy of structures, not in the simple

molecular components that are there to direct products of gene expression.” Life contains multitudes.

It is right to be amazed that it works at all. If, like Bill Clinton, you believe that credit for life belongs to God, I hope you might feel that They emerge looking far smarter and more inventive than the message of the Human Genome Project implied. If you don't feel a need to find a place for God, then I encourage you simply to allow yourself to be enchanted by the genius of life.

## **Fixing a Living Radio**

How we go about solving a problem reveals a lot about the nature of problem we consider it to be. In 2002, biologist Yuri Lazebnik, then at Cold Spring Harbor Laboratory in New York, found a memorable way to illustrate how we typically study biology.

He recounted how, as an assistant professor, he sought advice from a senior colleague about the perplexing whirlwind of activity taking place in his field (the study of spontaneous cell death, or apoptosis). What happens in biology, he was told, is that researchers beaver away in their recondite corners until some unexpected observation makes many think that what was previously a mystery may be soluble after all—and what's more, that this effort may result in a miracle drug. But as the topic booms and publications multiply in their hundreds or thousands, discrepancies and contradictions begin to appear, predictions fail, the problem looks harder than ever, and those drugs never materialize.

The generality of this scenario, wrote Lazebnik, “suggested some common fundamental flaw of how biologists approach problems.” To try to understand what that was, he followed the advice of one of his high-school teachers by testing that approach on a problem with a known solution. He set out to see if the methodology generally used in biology would work to show how a transistor radio works. How would that approach generally go? First, he wrote, researchers would persuade funders to let them buy a stack of radios that all work the same way, which they will dissect and compare with the broken one:

We would eventually find how to open the radios and will find objects of various shape, color, and size. We would describe and classify them into families according to their appearance. We would describe a family of square metal objects, a family of round brightly colored objects with two legs, round-shaped objects with three legs and so on. Because the objects would vary in color, we would investigate whether changing the colors affects the radio's performance. Although changing the colors would have only attenuating effects (the music is still playing but a trained ear of some can discern some distortion) this approach will produce many publications and result in a lively debate.

Another approach would be to remove components one at a time. Occasionally, some lucky researcher will find a part whose removal stops the device working at all. "The jubilant fellow will name the wire Serendipitously Recovered Component (Src) and then find that Src is required because it is the only link between a long extendable object and the rest of the radio.<sup>2</sup> The object will be appropriately named the Most Important Component (Mic) of the radio." And so on. Eventually, said Lazebnik, "all components will be cataloged, connections between them will be described, and the consequences of removing each component or their combinations will be documented."

Only then will the crucial question have to be asked: "Can the information that we accumulated help us to repair the radio?" And can it? In rare lucky cases, a fix might work—but the biologists won't really know why. Mostly, it won't work at all.

So what's wrong here? Lazebnik argued that biology is using the wrong language—a qualitative and sometimes personalized picture of "this component speaks to that one," rather than the true circuit diagram of an electrical engineer. Lazebnik's somewhat tongue-in-cheek paper made an extremely pertinent observation: the *modus operandi* of much of experimental biology might not be the one that will furnish a genuine understanding of how these systems work. Still, his prescription for doing better by developing a formalized engineering-style language was predicated on the analogy between a living system and a radio. He anticipated the objection that "engineering approaches are not applicable to

cells because these little wonders are fundamentally different from objects studied by engineers.” But he felt this was akin to a belief in vitalism.

That objection does not, however, follow at all. What if instead a radio simply is not the right analogy—if biology doesn’t work like any engineered system we have ever created? What if its *operational logic* is fundamentally different? Then we will need something more than a better formal language. We will need a new way of thinking—albeit not one that need invoke any mysterious vital force. I believe that this is the situation we face, and that both the successes and the failures of much biological research in the past two or three decades point to this conclusion.

In 2000 cell biologists Marc Kirschner, John Gerhart, and Tim Mitchison made a tongue-in-cheek allusion to vitalism in calling for a better way to understand life than by a detailed characterization of its parts and of their modes of interconnection. They “light-heartedly” called such an improved view “molecular vitalism,” saying,

At the turn of the twenty-first century, we take one last wistful look at vitalism, only to underscore our need ultimately to move beyond the genomic analysis of protein and RNA components of the cell (which will soon become a thing of the past) and to turn to an investigation of the “vitalistic” properties of molecular, cellular, and organismal function.

In other words, we don’t need some tautological “life force,” but we do need to ask what it is that distinguishes life from the lifelessness of its components. Only then will we have much hope of truly being able to fix a “living radio.”

To keep life running, we have to do a lot of fixing. The body goes wrong often, mostly in small ways but sometimes in big ones. We have become fairly adept at the mending process we call medicine, but often by trial and error, because we didn’t have good manuals to work from, but only occasional glimpses of how this part or that functions.

Already the emerging new view of how life operates within us is prompting some rethinking of medicine—of how we design drugs, say, and why some diseases such as cancer are so hard to prevent or cure. Some researchers now suspect that it might be time to shift the entire philosophy

underpinning medical research: for example, not to study and attack diseases one at a time, or to try to kill pathogens (that are typically smarter than us, adapting faster than we can retool our therapies) with bespoke magic bullets, but to take a unified view of disease. Many diseases wreak their effects through the same channels, and strategies for combating diverse diseases might involve similar or even the same approaches, especially involving the immune system.

And as we become more knowledgeable about where and when to intervene in life's processes, we can start to think of life itself as something that can be redesigned. Efforts to do so systematically began with genetic engineering in the 1970s, but that typically only worked well for the simplest forms of life, such as bacteria. What's more, it was limited by intervening only at one level of life's hierarchy: genetics. It was by no means clear that every desirable goal could be attained by tinkering with genes, and we can now see why: because genes don't generally specify unique outcomes at the level of cells and organisms.

Today we are beginning to redesign and reconfigure living entities, tissues, and organisms at several levels. We can reprogram cells to carry out new tasks and grow into new structures. We can create what some are calling multicellular engineered living systems: not mere blobs of living matter fed by nutrients in a petri dish, but entities with structure, form, and function, such as "organoids" that resemble miniature organs. Yet we are still very much in the foothills of this enterprise, trying to discern the rules that dictate the forms into which cells organize themselves. As our knowledge and our techniques improve, our ability to guide and select the outcomes becomes ever more profound. Some researchers believe that ultimately this will enable us to regenerate limbs and organs, and perhaps even to create new life forms that evolution has never imagined.

## **A Glimpse Ahead**

There's a lot in this user's guide because there is an awful lot to life. Modern biology is notoriously intricate, overburdened with fine details, arcane terminology, and impenetrable acronyms, and bedeviled by caveats

and exceptions that make it nigh impossible to make any statement without qualifications and footnotes.

It's my contention, however, that there is not *just* a lot to life. A common response to any attempt at generalization in biology is to say "Ah, but what about exception X?," almost as if it were a solecism to try to glimpse beyond all the trees to get a view of the wood (or the forest, if you are in the United States). Yet it is surely not the case that life is just a dizzying mess of fine details in which every aspect matters as much as any other. That can't be true, because no highly complex system can work that way. If this were how organisms are, they would fail all the time: they would be utterly fragile in the face of life's vicissitudes. It would be like making a mechanism from a billion little interlocking cogs in which, if just one of them snaps or jams or falls out of place, the whole thing will grind to a halt—and then expecting this machine to work for eighty years or so while being constantly shaken vigorously.

No, there are sure to be high-level rules that govern life, which do not rely on the perfect integrity and precise placement of all its parts. But if they are not summed up in the idea that we are "machines made [and defined and governed] by genes," then what are they?

It's a curious paradox that, while in recent years these principles have been becoming increasingly apparent, at the same time they have tended to be obscured beneath an avalanche of *data*. Data can be very valuable, indeed essential, for discerning general rules and patterns, but only so long as we do not end up fetishizing the data themselves (by literally making books from them, for example).

We have become extremely adept at gathering biological data, especially about the sequences of genomes, the structures of proteins and other biomolecules, and the variety of molecular components in cells and the interactions between them. By analogy with the science of genomics, these data sets are typically suffixed as "-omes": there are proteomes, connectomes, microbiomes, transcriptomes, metabolomes, and so forth. Thanks increasingly to the assistance provided by artificial intelligence and machine-learning algorithms, which can analyze far bigger data sets than humans can, we are able to survey and mine these -omes to glimpse the regularities and correlations within them. All this is immensely valuable, but in the end what it tends to offer are descriptions, not explanations. One

sometimes senses that some biologists prefer it that way—that they hope data mining will suffice for making predictions, so that we don't actually have to *make sense* of all the data or find coherent stories to tell about it. Instead, we can just rely on computers to find correlations between this data bank and that one. It's not clear, however, that this alone will enable us to make more and better interventions for human health. It's even less clear that it will act as a satisfying intellectual substitute for really understanding how life works.

With this in mind, I want briefly to suggest some of the themes and principles that will appear repeatedly in what follows, and which I hope might offer some common threads that can guide us through the challenging landscape.

*Complexity and Redundancy:* I once heard *Nature's* former biology editor say very wisely that in biology the answer is always “yes.” (One might argue that it is in fact “yes, but. . . .”) By this she meant that there are many different ways that a process can happen—that a signal can be transmitted within a cell, that a gene can be switched on or off, that cells can assemble into a particular structure. Traditionally this feature has often been regarded as a kind of fail-safe mechanism: because interactions between one molecule and another can't always be guaranteed to happen, evolution has provided backups. But in fact we'll see that the logic of biological redundancy is often of a different kind: there is a fuzziness to the system, so that different combinations of interactions can have the same result, and a particular combination can have different outcomes depending on the context. This, it seems, is a better way to get things done in a microworld beset by randomness, noise, and chance fluctuations.

*Modularity:* Life never has to start from scratch. Evolution works with what is already there, even if this means redirecting it to new ends. We might (with great caution!) compare it to an electronic engineer who uses preexisting circuit components like diodes and resistors, and standard circuit elements such as oscillators and memory units, to create new devices. Thus life possesses a modular structure. This is most obvious in the way large

organisms like us are assemblies of cells, as well as sharing common structures such as hearts and eyes. Modularity is an efficient way to build, since it relies on components that have already been tried and tested and permits the modification or replacement of one part more or less independently from the others.

*Robustness:* Life's resilience is remarkable. After a summer of terrible drought that saw all of England turn yellow-brown, it has taken only a few heavy rain showers for the green to start reappearing. Life is not invulnerable, but it is extraordinarily good at finding ways through adversity (which the world supplies in dismaying abundance). We will never have adequately explained life until we can understand where its robustness comes from. No doubt the aforementioned redundancy is a part of that, but robustness features in many contexts: in the way most embryos grow into the "right" shape, wounds heal, infections are suppressed, and more broadly, life on Earth has sustained its continuity for close to four billion years.

*Canalization:* Life is what physicists might call a "high-dimensional system," which is their fancy way of saying that there's a lot going on. In just a single cell, the number of possible interactions between different molecules is astronomical—and there are around 37 trillion cells in our bodies. Such a system can only hope to be stable if, out of all this complexity, only a limited number of collective ways of being may emerge. The number of possible distinct states that our cells adopt is far, far smaller than the number of ways one cell could conceivably differ in detail from another. Likewise, there are only a limited number of tissues and body shapes that may emerge from the development of an embryo. In 1942 the biologist Conrad Waddington called this drastic narrowing of outcomes *canalization*. The organism may switch between a small number of well-defined possible states, but can't exist in arbitrary states in between them, rather as a ball in a rugged landscape must roll to the bottom of one valley or another. We'll see that this is true also of health and disease: there are many causes of illness, but their manifestations at the physiological and symptomatic levels are often strikingly similar.

*Multilevel, multidirectional, and hierarchical organization:* To understand how life works, there is no single place to look. You will never find all the explanations at (speaking both metaphorically and literally) a single level of magnification. What is more, each level in the hierarchy of life's organization has its own rules, which are not sensitive to the fine details of those below. They have a kind of autonomy.<sup>3</sup> At the same time, influences can propagate through these levels in both directions: changes in the activity of genes can affect the behaviors of whole cells and organisms, and vice versa.

*Combinatorial logic:* It has been estimated that humans can discriminate between around one trillion odors. Quite what that number means is open to debate, but it is clearly very much larger than the mere four hundred different "receptor" molecules in our olfactory system: there is evidently not a separate molecular detector for each smell. The different odor sensations must arise from different patterns of activation of this relatively small set of receptors. That is, the smell signals our brains receive are *combinatorial*. Think, for comparison, of how just three light sources (red, green, and blue-violet) in visual display screens can create a whole gamut of colors through differences in their relative brightness. Molecular signals that are combinatorial, rather than relying on unique molecules to supply different outputs, are widely used in biology, probably because they are economical in component parts, versatile, adaptable, and insensitive to random noise: all of them attributes that serve life well.

*Self-organization in dynamic landscapes:* Many things are possible in life, but not everything. Evolution does not select from an infinite palette: there are specific patterns and shapes in space and time that arise out of the complex and dynamic interactions between the components of biological systems, much as there are common features of cities or animal communities, or of crystal structures or galaxies. Think of it rather like rain falling on a landscape: the water itself is not programmed to flow in any particular direction, but the shape of the landscape causes it to gather in some places and to move away from others. The

language of landscapes, basins, and channels is often useful in biology.

*Agency and purpose:* Agency is becoming something of a buzzword in some biological circles, especially those concerned with processes of cognition. The trouble is, no one seems able to agree on what it means. Intuitively, we might suspect that what distinguishes living organisms from nonliving matter is this notion of agency: they can manipulate their environments, and themselves, to achieve some goal. This makes agency inextricably linked to ideas about purpose. That is probably why the problem of agency has been (absurdly) neglected for so long in the life sciences, where questions of purpose have long been shunned as quasi-mystical teleology, perhaps only one step away from the dreaded concept of intelligent design. The result of this neglect and avoidance is that we can end up skirting around the most characteristic feature of all life. I propose that the time has come to embrace it—and that there is nothing to fear in doing so.

*Causal power:* One of the biggest obstacles to understanding how life really works has been a failure to get to grips with causation. It's a hard problem, not least because causation is a vexed topic in its own right; philosophers still argue about it. We already know from daily experience how difficult it is to decide what counts as a cause of a phenomenon. Are the words appearing on my screen being caused by the impacts of my fingers on the keyboard, by electrical pulses within my computer's silicon chips, or by the more abstract agency of my thoughts and feelings? But these questions are not intractable, and we do have some conceptual and mathematical tools for handling them. Too often, causation in biology, as indeed in the world in general, has been assumed to start "at the bottom" and filter up—so that, for instance, characteristics at the level of an organism's traits are deemed to be "caused" by genes. As we'll see, we can gain a better understanding of how life works, and how to intervene in it effectively, when we take a more sophisticated view of biological causation.

If everything in this book is correct, it will be a lucky miracle, and no reflection on my depth of understanding or intellectual powers. I suppose that is hardly a statement to inspire great confidence in what you are going to read, but the honest truth is that I am writing about issues that are still being debated by experts, sometimes with vehemence. Nevertheless, I believe there is no serious doubt that the narrative we ought to be telling about how life works has shifted over the past several decades, and it is time we said so. Given how increasingly important the life sciences—from genomics to precision medicines and research on aging, fertility, neuroscience, and more—are becoming in our lives, I believe this is nothing less than a duty. The historian of science Greg Radick has argued that we should “teach students the biology of their time,” and not the tidy simplifications concocted a half-century or more ago. He is right—but we should teach it to *everyone*.

The new story that is emerging is, it’s true, sometimes more complicated than the old half-truths. But I think this story is coherent, cogent, and consistently supported by many independent strands of research in genetics and molecular biology, cell biology and biotechnology, evolutionary theory, and medicine. Many of the details remain unclear and contentious, but the broad outline seems now unassailable and, I believe, exhilarating in what it tells us about the astonishing process that created a form of matter able to begin understanding itself: us. What’s more, this new view of life plugs us back into the universe. It does not replace or undo older ideas about natural selection but deepens them to help us see what is truly different and special about living organisms: what it really means to be alive.

# The End of the Machine

## A New View of Life

Marjorie, then eighty-eight years old and living in a nursing home, was among the millions of people infected with the coronavirus during the COVID-19 pandemic that began in 2020. She was frail and asthmatic, and she suffered from the inflammatory lung disease COPD. “If I get it, I’m finished,” she had told me before her infection.

Another person who caught COVID was Ray, a fifty-six-year-old man in good health and with no previous complications that would put him on the danger list.

One of these two people—they are both real, but I’ve changed their names—tragically died from the effects of the virus. And of course I would not be setting up the situation in this way if it had gone in the direction you would predict. No one was more astonished than Marjorie when she made a quick recovery from the virus.

There are countless stories of this sort: of sad, unexpected deaths and of unlikely escapes. While it was clear that older people were statistically at greatest risk from COVID-19, no one knew quite how their own body would respond to infection. Many people had the virus without even knowing it, quite possibly transmitting it unawares to others who would die from its effects. The vast majority of those infected did not die, but many developed serious and long-term health problems of bewildering variety, ranging from brain damage to blood clotting, persistent exhaustion to heart problems.

The pandemic reminded us in a terrible manner how little we understand about our bodies and about how they are assailed by the slings and arrows of outrageous fortune. And yet in one sense we knew, right from the outset, everything about the SARS-CoV-2 virus responsible for it all. No sooner had the virus been isolated when it first emerged in Wuhan, China, than its

genome—a relatively short stretch of RNA (for the coronavirus, like many other viruses, encodes its genes in the RNA molecule, not in the closely related DNA that is the genetic fabric for all cellular organisms from bacteria to us)—was sequenced and the protein molecules it encodes were characterized. We quickly discovered the molecular-scale details of how the virus attacks and enters human cells: the so-called spike protein on its surface latches onto a protein called ACE2 on human cell surfaces.

The hard part was to understand what happened next. Sometimes the virus might send the body of an infected person into a kind of immune overdrive, damaging their lungs and their ability to absorb oxygen. Sometimes, on the contrary, the infection produced no symptoms at all. One of the (many) reasons why the controversial idea of “focused protection” as a pandemic strategy—sheltering the vulnerable while allowing the virus to infect those unlikely to greatly suffer from it, until herd immunity was attained—made no sense is that we had no idea, other than via crude statistical demographics of age and preexisting health conditions, who the “vulnerable” actually were.

Despite this lack of understanding, we were able to develop vaccines in record-breaking time that have done an excellent job of protecting most people from the worst ravage of the virus. We knew how to use harmless protein fragments of the virus, or pieces of RNA encoding them, to stimulate our bodies’ immune defenses, triggering them to produce antibodies that attach to the virus and block its action or flag it for destruction.

Here too, though, the consequences were unpredictable. Most people who had two doses of a COVID-19 vaccine only became mildly ill if infected. (Why were two needed, and not just one, or ten? We don’t yet really know.) But a small proportion of unlucky individuals got seriously ill or even died from COVID-19 despite being vaccinated. Meanwhile, among the millions of people who took the vaccine, the vast majority merely felt tired or ill for a day or so, as if with a mild case of flu. Many noticed no side effects at all. But a tiny minority suffered unpleasant side effects, especially blood clots that could be life-threatening. The chances of this were minuscule—much smaller than the chances of nasty consequences if you caught the virus without being vaccinated—but still you mostly just

had to hope that you weren't one of those very few who drew the short straw.

This is surely a curious combination of circumstances. We have mighty technologies for characterizing our pathogenic foes and for developing medicines against them. The COVID vaccines, especially in the rapidity of their creation and testing, have been one of the greatest triumphs of modern science. And yet in some ways we seem little better off than we were in the Middle Ages, seeking medicines (including COVID antivirals) largely by trial and error, and having to hope that, if we're infected, our god or blind luck will spare us. How can this be? Why can't we do better? If we can "decode life" down to the atomic scale, what are we still missing?

## **A Brief History of Life**

In ancient times, people didn't particularly look for metaphors to understand life. More often, they used life itself as a metaphor to understand the world. Life seemed to be the organizing principle of the cosmos.

But as for what it *is*—that was almost like asking what the classical elements (air, water, and so on) were. Life was a fundamental property, not something that could be decomposed into ingredients. For Aristotle, the aliveness of living things was imbued by their soul (*psyche*). This is not to be confused, although later it would be, with the Christian notion of a soul; rather, it refers to a kind of innate capacity for action. The *psyche* had no substance in itself, but it was inseparable from the body: it was in the very nature of living bodies. Aristotle believed that a living body's soul gives it various capabilities for growth and self-nourishment, movement and perception, and intellect. Different kinds of living bodies have different degrees of soul: plants are capable only of growth and nutrition (they have a vegetative soul), animals may also move and have sensation (a sensitive soul), but only humans have the rational soul that also conveys intellect.

With the rise of a mechanistic view of the world in the seventeenth century—the idea that all of nature can be understood on the basis of forces acting between particles in motion—life became conceptualized as a kind of machine. The mechanistic philosophy reached its apotheosis with Isaac

Newton's laws of motion, laid out formally in his epic 1687 tract *Philosophiae Naturalis Principia Mathematica*, but this vision of a machine-cosmos was already well-established by then. In his *Discourse on the Method* (1637), René Descartes set out a view of the human body as a wondrous mechanism of pumps, bellows, levers, and cables. All of these parts are animated by the divinely granted rational soul, which is lodged in the body but, contra Aristotle, not dependent on that physical host (for it would have been heresy to deny the immortality of the soul). Descartes set out this mechanistic vision of the human body most extensively in his *Treatise on Man*, which he began in the 1630s but abandoned when he witnessed the consequences for Galileo of advocating philosophical ideas that might be considered to conflict with holy scripture. (The *Treatise* was published posthumously in 1662.)

The mechanistic picture of living things was taken further by the French physician Julien Offray de La Mettrie, whose *Natural History of the Soul* (1745) seemed to deny the need for that notion at all. Life was an innate property of the living body, he said, not some supernatural force that sets the parts in motion. As he wrote later, the human body is a “machine which winds its own springs.” To the extent that we have a soul at all, it is a kind of emergent property of our complexity of organization, the summed complement of a fundamental “irritability” of the fibrous tissues of the body. The book was denounced as blasphemous and La Mettrie had to flee from Paris to Leiden, where in 1747 he published an even more trenchant defense of the mechanical view of life, *L’homme machine* (*Man, a Machine*). Here he presented humankind as no different from the “perpendicularly crawling machines” that are beasts. All that distinguishes us, he said, is a great complexity in the arrangement of our irritable fibers.

La Mettrie's books got him into trouble, but by this stage of the Enlightenment the church was fighting a rearguard action against the increasing authority of science to speak to the nature of organic, living matter. By the late eighteenth century, chemists such as Antoine Lavoisier in France were analyzing living matter in the literal sense: breaking it down into its constituent elements and studying the chemical principles, such as respiration, on which it depended.

All the same, it remained profoundly puzzling what distinguished a carbon-based organism from a piece of diamond, given that both could be

combusted into (as we'd now see it) carbon dioxide gas. Some suspected that the difference was merely material: there was some special form of substance that was inherently alive by virtue of its chemical composition. The French naturalist George-Louis Leclerc, Comte de Buffon, postulated a kind of matter called *matière vive*, composed of “active molecules” with an innate tendency to move—a kind of “little life” that is “primitive and apparently indestructible.” The life of organisms is then just the result of “all the actions, all the separate little lives.” These living molecules also possess a kind of primitive intellect from which that of animals arises<sup>1</sup>. This “atomized” view of life as the sum of its molecular parts was shared by the great systematizer of the Enlightenment, Denis Diderot, who speculated about how a swarm of such “living points” can create “a sort of unity which exists only in an animal.” Thus life arises from a kind of “vital force” that animates its ingredients.

Buffon's notion of a kind of primitive “living matter” was shared in the late eighteenth century by the Scottish surgeon John Hunter, who dignified it with the Latin term *materia vitae* without thereby shedding any new light on what it might be. But in 1835 the French anatomist Felix Dujardin claimed to have identified something of the kind: a gelatinous substance made by crushing microscopic animals, which he named *sarcode*. It was subsequently renamed *protoplasm*, and Austrian biologist Franz Unger suggested that it might be a form of the organic substance called “protein,” which was then recognized only as a nitrogen-rich organic material common in living things. In the 1850s the English zoologist Thomas Henry Huxley claimed to have isolated protoplasm—the “physical basis of life”—from sediments dredged up from the sea floor, which contained carbon, nitrogen, oxygen, and hydrogen. Its living character, he said rather vaguely, resulted from “the nature and disposition of its molecules.” In fact Huxley's protoplasm turned out, to his chagrin, to be nothing more than a gel produced by chemical reaction between seawater and the alcohol used as a preservative for the organic matter in the sediments.

The idea of a “vital force” was hardly an answer to the puzzle of life. It was, rather, a tautology that just displaced the question: things are alive because their component parts are. Whence does the vital force arise? In the early nineteenth century, some scientists suspected it might be of an electrical nature, given how electricity discharged from storage devices known as Leyden jars could make the dissected limbs or dead bodies of

animals twitch with apparent animation. At any rate, by demonstrating a continuity between the chemical composition of “organic” substances derived from living organisms and inorganic substances made from evidently inert matter such as salts and gases, nineteenth century chemists eroded the idea of a distinct form of matter that is inherently alive.

In 1812 the great Swedish chemist Jöns Jakob Berzelius dispelled the idea that life could be explained by some mysterious vitality inherent in matter by virtue of its composition. “The constituent parts of the animal body,” he wrote, “are altogether the same as those found in unorganized matter, and they return to the original unorganic state by degrees . . . after death.” He despaired of getting to the bottom of the mystery, saying that “the cause of most of the phenomena within the Animal Body lies so deeply hidden from our view, that it certainly will never be found.” In seeking for it, he said, “the chain of our experience must *always* end in something inconceivable; unfortunately, this *inconceivable something* acts as the principal part in Animal Chemistry.”

All the same, Berzelius added a fruitful notion. Rather than postulate some “vital force”—“a *word* to which we can affix no idea”—we should recognize that “this *power to live* belongs not to the constituent parts of our bodies, nor does it belong to them as an instrument, neither is it a simple power; but the result of the mutual operation of the instruments and rudiments on one another.” In other words, it is not so much a question of what the molecules *are*, but of what they *do*, and specifically, of what they do collectively.

To that degree, then, life becomes a question of how its components are *organized*. The question of organization came increasingly into focus over the course of the nineteenth century as microscopic methods improved to the point that researchers could look at living things below the level of the cell.<sup>2</sup> That all life is cellular was proposed in the 1830s by the German zoologist Theodor Schwann, who wrote in 1839 that “there is one universal principle of development for the elementary parts of organisms, and this principle is in the formation of cells.”<sup>3</sup> Schwann’s colleague, botanist Matthias Jakob Schleiden (the two worked in the Berlin lab of physiologist Johannes Müller), believed that cells were spontaneously generated within organisms, but another of Müller’s students, Robert Remak, showed that cells multiply by dividing. That notion was popularized and extended by yet

another Müller protégé, Rudolf Virchow, who coined the memorable phrase (if your Latin was up to scratch) *omnis cellula e cellula*: all cells come from cells. For Virchow, complex tissues and organisms are collectives of this fundamental unit of life, which is a kind of “elementary organism” in its own right.

Toward the end of that century, microscopic studies of cells showed that they were no mere blobs of protoplasm-like matter but had internal organization of some sort, visible as dark blobs, fibers, and other structures that could be rendered more apparent by using dyes to stain them. There were little granules that were named “mitochondria” in 1898, spongelike membranes, and fibrous bodies labeled “chromosomes” (“colored bodies,” referring to their ability to be stained by dyes). It wasn’t clear what all this internal organization was for, but it showed that cells have components and compartments of some kind, and an understanding of how they work would surely demand that we characterize these structures in more detail.

That was hard—because they were so small, so numerous, and so varied. Cell biologists could see changes occurring in the internal organization as cells went through their cycle of repeated division. But understanding the causes and significance of these transformations was another matter. All we can do, said French physiologist Claude Bernard in 1878, is to “observ[e] the facts nearest to us, [and] advance step by step till we finally reach the determinism of these fundamental phenomena.”

But piling up facts won’t do; we need to understand general principles. In the early twentieth century, the word *organization* was thrown around as a kind of catch-all invocation of aspects of life barely understood even in broad outline. “We are forever conjuring with the word ‘organization’ as a name for that which constitutes the integrating and unifying principles in vital processes,” admitted the American cell biologist Edmund Beecher Wilson in 1923. This is a common pattern in biology, which began with terms like *soul*, *vital force*, and *protoplasm* and, as we’ll see, has continued by referring to such concepts as *gene action* and *regulation*: terms that label things and processes barely grasped. This is not a failing of the science, however, but a necessary tool for dealing with life’s dizzying complexity. It’s better to have a vague concept that may act as a bridge across a void of ignorance than to come dejectedly to a halt at the brink.

## The Value and Dangers of Metaphor

There were, and still are, many disparate fronts on which scientists try to understand how life works. Some study it at the scale of the cell, characterizing all those exotically named components and their functions: the nucleus, the mitochondria, the Golgi apparatus, endoplasmic reticulum. Developmental biologists, meanwhile, try to figure out how cells grow, specialize, and create tissues with particular shapes and locations in the progression from fertilized egg to embryo to organism. And as some biologists wrestled with the cell's "organization" in the early twentieth century, others were trying to understand the principles of heredity and how these were connected to the entities that had been christened genes—as well as how those processes related to the "great chain of being" in Darwin's theory of evolution by natural selection. And still others pursued the chemists' perspective on life by looking at its molecular nature, in particular the biochemical transformations involved in metabolism and the role and nature of the molecules called enzymes, made of protein, that acted as catalysts for those reactions. Each of these pursuits was and is immensely difficult and demands a deep stock of specialized knowledge, such that biologists working in one field may find that they scarcely share the same lexicon—or worse, that they use the same words for different purposes. They do not necessarily concur about which are the most important questions to ask about how life works.

What they do all share in common, however, is a strong reliance on metaphor. To some extent that is true of all science—indeed, of all language, even all thought. But biology perhaps has greater need of it than other sciences precisely because the principles seem so hard to grasp and to articulate. Favored metaphors change over time, but—and this is less often appreciated—that does not simply mean that one supplants the other. The concept of "vitalism" might be traced back to the Aristotelian soul and is generally regarded as obsolete in biology today, but in fact we'll see that it still survives in cryptic forms, most particularly in the way biomolecules and other reductive components of life may be unconsciously attributed a kind of agency they do not really possess. The Cartesian mechanistic metaphor is very much alive and well: biologists routinely speak of "molecular machines" such as enzymes, and not without good reason. But such language can morph into a literal view, in which we might really treat

microscopic biological entities as though they were cogs and motors that operate in the same way as our technological ones. This, as we'll see, can be deeply misleading.

The organizational metaphor, meanwhile, is apparent in the way cells are commonly described as tiny “factories,” within which biomolecules are the workers that collaborate to churn out exquisitely crafted molecular products, using energy from the “powerhouses” of the mitochondria and creating waste that must be disposed of or recycled. To these older metaphors was added another in the second half of the twentieth century: *information*. In the age of the digital computer, biologists became convinced that life itself was a kind of computation, an algorithm dictated by a digital code of instructions imprinted in the storage tape of DNA. “Today,” said François Jacob in 1970, “living organisms are seen as the site of a triple flow of matter, energy, and information.” Rarely, he claimed, had a metaphor imposed by a particular technological epoch been more apt.

All of these metaphors have their uses, for they were not coined without good reason—and I will sometimes draw on them. But the old saw that the price of metaphor is eternal vigilance<sup>4</sup> is nowhere more apt than in trying to understand life. Metaphors in biology have a dangerous tendency to turn into “explanations,” and schematic representations of experimental findings—how a set of molecules appear to interact, say—may be mistaken even by experts for literal pictures of what happens.<sup>5</sup> One of the fundamental messages of this book is that we cannot properly understand how life works through analogies or metaphorical comparison with any technology that humans have ever invented (so far). Such analogies may provide a foothold for our understanding, but in the end they will fall short, and will constrain and even mislead us if we don't recognize when to relinquish them. To truly understand, say, how an embryo grows, says developmental biologist Jamie Davies, “we must be prepared to move beyond homely analogies, based on how we build things, and see the embryo in its own terms.” One obvious but very profound distinction of life from machines, for example, is that life must be sustained continuously or not at all: you can't turn it off and on again.

Comparing life to a machine, a robot, a computer, sells it short. Life is hard to understand precisely because it is like none of these things we have created. And when we forget the limitations of metaphor, our science fails


and its applications flounder. We might seek inappropriate or ineffective medical interventions, for example. The fundamental problem with the machine metaphor of life—which applies equally to the “electronic circuit” metaphor—is that it compels us to consider the parts as *things with functions*. We take up a particular component and ask: what role is *this* part playing?

But as we will see, there are constituent elements of living things (such as proteins) for which that might not always be a meaningful question. It’s not just that the role of an entity depends on its context (which can be true of machine parts too); the whole concept of a “role,” and of a *mechanism* in which the role is enacted, becomes murky. Biologist Robert Rosen summarized this conventional approach of reducing life to its parts nicely in 1991 when he said it amounts to the notion that we can (and indeed should) “throw away the organization and keep the underlying matter” (see [box 1.1](#)). Rosen complained that molecular biology seemed as a result to have relinquished any attempt to explain, and was content instead to merely describe, or stand and watch. Even what it describes, he said, tends to be weirdly static, idealized, and disembodied: there is no movement or dynamics, no noise, and very little sense of things being organized in space. The machine metaphor, Rosen claimed, had become so ingrained that to question it was anathema, perhaps of a suspiciously mystical kind. “To suggest otherwise,” he wrote, “is regarded as unscientific and viewed with the greatest hostility as an attempt to take biology back to metaphysics.” It’s time to let go of such prejudices.

## **Are Machines Even “Machines”?**

Perhaps this also means letting go of prejudices about machines themselves. For the fact is that even machines are not what they used to be. When the metaphor is used in biology, typically it is intended to conjure up images of mechanisms or robots with moving parts that respond in well-defined, deterministic, and transparent ways to input signals: at its simplest, systems of cogs and levers. That much is badly misleading. More recently the analogy has been made with electronic circuits, where, for example, electrical signals are directed along channels by switches and junctions

carefully designed and laid out by electronic engineers. That too is a poor metaphor. So long as we insist that cells are computers and genes are their code, that proteins are machines and organelles are factories, the picture that emerges is a clumsy marriage of the mechanical and the anthropomorphic. Life becomes an informational process sprinkled with invisible magic.



**Fig. 1.1** The structure of the bacterial flagellar motor, deduced by researchers at Zhejiang University in China using the technique of cryo-electron microscopy (see p. 151). Image courtesy of Xing Zhang and Yongqun Zhu, Zhejiang University; see Tan et al. (2021).

Now, I would certainly admit that even the traditional view of the machine seems apt for *some* biological entities. Take the so-called bacterial flagellar motor, an assembly of protein molecules that sits in the cell membrane of bacteria and enables its whiplike flagella to spin and propel the organism through water in a corkscrew-like fashion (fig. 1.1). Molecular assemblies like this one rightly leave us awed about the kinds of structures and devices nature can generate. They also help to explain why proteins are commonly described as molecular machines. What could be more reminiscent of our own technological devices than this—an axle built to rotate inside a confining sleeve?

The machine metaphor for biomolecules was promoted by biologist Bruce Alberts in an influential 1998 article titled “The Cell as a Collection of Protein Machines.” The entire cell, he wrote, “can be viewed as a factory that contains an elaborate network of interlocking assembly lines, each of which is composed of a set of large protein machines.” The truth is, however, that rather few biomolecular structures are as seemingly literal as the flagellar motor in translating familiar mechanical notions—here, of a rotary motor—to the microscale. More to the point, we should not expect the principles of our own machinery to translate in a straightforward way to the molecular scale, where the roles and even the nature of phenomena such as viscosity, friction, rigidity, and adhesion are very different. As philosopher of science Daniel Nicholson has said,

Owing to their minuscule size, cells and their macromolecular components are subject to drastically different physical conditions compared with macroscopic physical objects like machines, and . . . using machine metaphors to explain microscopic phenomena is consequently more likely to obscure and deceive than it is to elucidate and enlighten.

I think Nicholson is right. As we’ll see, it is not just that molecular-scale phenomena such as molecular vibrations and randomness make the machine metaphor a little fuzzy; in general, proteins, and Alberts’s “assembly lines,” don’t employ the same principles at all. And living organisms as a whole do not work like machines that resemble any we have ever made.

Biologist Michael Levin and computer scientist Josh Bongard argue that we might turn this fact on its head and reconsider our notion of what a machine is: to regard living things as “machines as they could be.” “We view life as an especially interesting class of machines that is making us expand the limiting old ideas of what machines are and how to make them,” says Levin. We should not make life fit the image of our present-day machines, but we might yet make our machines in the image of life. Thus we might, in the future, use living systems as the inspiration for new kinds of artificial device. In a sense we have already begun to do this. Today’s information technologies are creating a hazier picture of design and function, for example by being tolerant of faults, noise, and errors in ways that don’t simply rely on redundancy: on having backup circuits to handle component failures, say. It’s becoming harder to say, or even to discern, just *how* the device works.

The artificial intelligence algorithms that, by analogy with the brain, we even call “neural networks” are a good example of this. These networks function not by design—doing their job the moment they are switched on—but through learning and training. They are wired somewhat like the tangles of interconnected neurons and synapses in our skulls, and we have treated them a little like black boxes: they work, but we’re not sure what rules they use. We have also developed computer algorithms that improve their own performance not by painstaking debugging but by a process that mimics evolution: making random changes and favoring ones that work. Functionality arrives not by design based on deep understanding but by trial and error combined with a selection procedure to assess the effects of such scattershot change. But increasingly, such bio-inspired approaches to design don’t seem to be enough. To make improvements and to have confidence in what emerges, we need to know something about *how* these systems work. I am confident that biology will point the way to such understanding.

Some might argue, then, that a “machine” is any entity that effects some change prompted or guided by environmental stimuli. Almost by definition, *that* type of “machine” seems a reasonable description of biological cells and organisms. My aim, though, is not to demolish the machine metaphor so much as to complicate it: to demand that, like all of the popular tropes of biology, it not be used lazily, excessively, or misleadingly. I am not sure biology is generally very vigilant about its metaphors.

## What Is Living All About?

There's another crazy question I must ask at this stage. It's crazy because, given that it has been debated furiously for millennia without any consensus, there is no prospect that I will answer it here. All the same, it's important to raise it at the outset.

Here we go, then: What even *is* life?

Commonly, the question has been interpreted as a demand for a set of criteria against which we can judge whether an entity is "living" or not: a kind of tickbox checklist. And so we see claims that, for example, life must be self-reproducing, or must undergo Darwinian evolution, or must embody complex self-organization, and so on.

Then problems arise because we find that some entities make the cut that we feel should not, and some do not that we feel ought to. This is probably an inevitable consequence of having only one set of interrelated living things to go by. Life on Earth is astonishingly and wonderfully diverse, from single-celled parasites to elephants, but it all came from the same source and so all shares, for example, the features of being cellular, being water-based, and possessing DNA and proteins and so on. We can't decide whether or by how much those shared attributes are essential or, on the contrary, just parochial (see [box 1.2](#)).<sup>6</sup>

Happily, I'm not obliged here to try to offer a definition of life, nor to consider the equally vexed issue of how it began. Eons of evolution insulate us now from the mysteries of life's origin—time enough for us now to treat it as a different *kind* of matter from the rocks and oceans of the early Earth, not in terms of its fundamental constituents but of how they got to be where and what they are. I'm not obliged to define it because my topic here is already laid out in plain view. We are ourselves arguably the most puzzling variety, not because we represent any pinnacle of evolution, but because we are able to ask how life works from the curious position of being inside it.

My subject, then, is how the living things *we know about*, and in particular, us and other complex creatures like us, actually work. Some of these workings of the human organism don't generalize to bacteria and other single-celled organisms, and so we certainly won't need to fret about whether, say, they may apply also to viruses, which occupy a disputed territory between the living and the inorganic.

All the same, I am not going to totally evade that question above, because the *how* of life can't wholly be disengaged from the *why*.

Already I'm in hot water. Science, we are often told, is not supposed to ask *why* questions, because that way lies teleology, or God (it is hard to know, for many scientists, which is worse). The more we understand about the universe, wrote physicist Steven Weinberg, "the more it also seems pointless."

But if it is true that the universe is pointless, why do so many people feel uncomfortable with, even antagonistic to, that suggestion? Surely because it conflicts with what we experience. The life of even the most atheistic and dispassionate of scientists is filled with moments that very much "have a point"—moments that *matter*. One couldn't sustain meaningful human relationships if that were not so. If life had no point, why would we bother sustaining it at all?

I'm quite certain that Weinberg, a deeply read and thoughtful humanist, would have concurred that his view doesn't preclude the existence of meaning for people. Of course we care about others and about ideals and principles, of course we seek and feel purpose. But in the grand scheme of things (so the Weinbergian argument would go), this is all very parochial. Sure, it matters a lot to me that I don't miss my flight, or that my mother is being well cared for in the hospital—but that can be hardly supposed to matter for the inhabitants of Trappist 6a, an Earth-like planet orbiting a star forty-one light years away,<sup>7</sup> or, for that matter, to anyone alive (if there is anyone alive) six generations in the future.

Weinberg's vision of a universe without purpose or meaning has become so much the scientific orthodoxy that it is almost obligatory for biologists to insist on it too. Words like *purpose*, *meaning*, even *function*, are treated with a caution bordering on disdain in the life sciences. At best they are corralled with scare quotes that proclaim them mere figures of speech; at worst they are excoriated as signifiers of a surreptitious religiosity.

Such aversion has often led biology to deny its own nature. For one of the best ways to characterize living entities is not via any specific features—replication or evolution, say—but the fact that for them, there is meaning. Things in their environment may take on meaning. Life, we might say, is that part of the universe that is not "pointless." And the fact that this "point"

is not merely parochial but in fact entirely personal is not to be sniffed at. On the contrary, *this is all it can mean*.

For I am afraid that Weinberg's much-quoted remark is, when considered carefully, itself without meaning. It makes a category error, using words where they do not belong. It is rather like saying that the more we understand about water, the more it seems friendless. There is no more reason why we should expect to find meaning in cosmology or particle physics than we should expect to find happiness or wisdom there. Meaning is not some mysterious force or fluid that pervades the vacuum. No; life is what *creates* such meaning as exists in the cosmos. Only for living things—or, to speak more generally, for things that, by their very nature, are imbued with purposes and goals—can there be a “point.” I suspect it is in fact precisely by virtue of being a thing that has autonomous goals, and that can autonomously attribute meaning, that an entity can be said to be alive.<sup>8</sup>

For some feature of an organism's environment to acquire meaning, the organism doesn't have to be “aware” of it. I don't think (although some biologists would dispute this) that a bacterium is conscious of its environment. The organism simply needs mechanisms for evaluating the value of that feature and acting accordingly. Looked at this way, life can be considered to be a *meaning generator*. Living things are, you could say, those entities capable of attributing value in their environment, and thereby finding a point to the universe.

Does this work as a definition of life that includes all we'd want it to and excludes the rest? Probably not. Could we make mindless machines that could mimic the meaning-generation of life? Probably—just as we can devise computer algorithms that mimic all kinds of lifelike features, such as self-replication and evolution. I don't think that's a problem. I'm not concerned with arbitrating what we admit to, and reject from, the Life Club. Rather, I want to introduce this idea because I think it needs to sit at the base of all consideration of how life works. A key reason why the machine metaphor for life is limited is that it doesn't include the possibility of meaning. To get to a machine that creates meaning, we need to move far beyond the traditional conception of a machine.

Similarly with the idea of life as computation. A computational algorithm that takes input data and generates some output from it doesn't really embody any notion of meaning either. Certainly, such a computation does

not generally have as its purpose its own survival and well-being. It does not, in general, assign value to the inputs.<sup>9</sup> Compare, for example, a computer algorithm with the waggle dance of the honeybee, by which means a foraging bee conveys to others in the hive information about the source of food (such as nectar) it has located. The “dance”—a series of stylized movements on the comb—shows the bees how far away the food is and in which direction. But this input does not simply program other bees to go out and look for it. Rather, they evaluate this information, comparing it with their own knowledge of the surroundings. Some bees might not bother to make the journey, deeming it not worthwhile. The input, such as it is, is processed in the light of the organism’s own internal states and history; there is nothing prescriptive about its effects.

There is currently no well-developed theory of meaning in science, nor are there theories for understanding the related concepts of purposes, goals, and intentions. In part, this is because the very existence of such factors in how life works has often been denied; at best, they tend to be regarded as “as if” properties, which it merely *looks like* living things possess. This is as peculiar as supposing that consciousness is merely an illusion that we believe we have—that stance denies the very phenomenon it is adopted to describe. In [chapter 9](#) I suggest how meaning, purpose, and goals can be made respectably scientific attributes.

The reason life can generate meaning (for itself) is that it *evolved*. It may be that we can imagine other kinds of system we would want to call life that have not evolved by natural selection, or indeed by any other means. I’m agnostic on that issue—but again I don’t need to wrangle about it here, because I’m talking about life that *did* evolve through Darwinian natural selection.

We have to tread carefully here. There is absolutely no reason to suggest that the *purpose* of evolution itself is to produce meaning-generators; evolution has no purpose that we can discern, nor any reason to have one. It’s the other way around: meaning-generators are successful entities in a Darwinian world. Making meaning is a great way of staying alive and propagating—so much so, indeed, that it’s probably the only way to be alive at all.

Any explanation of how life works must take account of its evolved nature. This in turn means that even if there is no reason to suppose that the

mechanisms of life are optimal, as good and effective as they could possibly be, it's fair to suppose that there is some advantage to them being the way they are and not otherwise. If we find (as we do!) that life doesn't work by carefully passing information from the genome up through tightly orchestrated supply chains of molecules until the blueprint has been realized in the organism, we might reasonably conclude the reason that this strategy is not used is that there are better ways of making an organism (and perhaps even that this strategy couldn't work at all).

These ideas may help to sharpen the notion of *function* that is commonly invoked in biology. To ascribe an entity a function is to suppose that it has goal-directedness. While biologists in the early twentieth century fixated on the concept of organization as a defining feature of life, in fact organization per se is by no means unique to living matter. Crystals are highly organized, at least in the sense that they are highly orderly. Chemists speak routinely of “self-organized” molecular systems that can spontaneously form complex structures without any involvement of life processes, purely as a result of the play of forces between the constituents. As we'll see, living systems make use of such self-organization because it is a cheap way of creating order and structure: it doesn't need lots of detailed encoding and guidance but is granted “for free.” But the organization we generally see in living systems is not of the sort that can appear in nonliving ones. The intricate structures that microscopists see inside cells, and even the molecular-scale organization of “devices” such as the protein-manufacturing ribosome or the light-harvesting complex of photosynthesis, have the form they do because they have a function. They have acquired it by means of natural selection, which is fundamentally a goal-creating process. Such language can make biologists nervous,<sup>10</sup> but it need not. It says only that evolution is a process by which goals and functions are created—a process that may arise spontaneously in nature and thus in some sense is inherent in physical laws.

Biologists often like to quote the title of a 1973 essay by Ukrainian American evolutionary biologist Theodosius Dobzhansky: nothing in biology makes sense except in the light of evolution. I'm not sure, though, that the deepest implication of Dobzhansky's statement is always understood. It is not simply saying that evolution is the universal process by which life has been molded—and it certainly cannot mean that as a consequence all we find in biology has been shaped and dictated by the

adaptationist requirements of Darwinian natural selection. Rather, the point is that we need to acknowledge what evolution does to matter: it gives matter goals and functions. That is what makes evolved life so special.

### **Box 1.1: How Much Does Reductionistic Data Collection Help Us Understand How Life Works?**

To judge from all the efforts life scientists make to gather immense data sets, you might suppose that their discipline is hampered by lack of information. Huge international projects exist (and have existed) to map out the entire genome, to characterize every protein and RNA molecule and every way in which they interact, to create maps of the brain or of developing embryos at the scale of individual cells. Sometimes it may seem that no sooner have we filled up one vast database than we move on to the next, before we have even interpreted what we have gathered already. Sometimes this rush to the next Big Data challenge is justified with the implication that *that's* what we needed all along. And yet all this information is sometimes gathered in the absence of what science really needs to make progress: hypotheses to test. It's almost as if there's a belief that insights will simply begin to seep out of the data bank once it reaches a critical mass.

These biological data sets can yield important insights, and I will draw on some of them in this book. What's more, plenty of biologists *are* deep, synoptic thinkers; most of what I will say is indebted to them. And biology has so much data to collect because life is so very complicated, so heavily populated with diverse component parts. But there are two related problems for a mindset that demands more data while remaining skeptical of theories to explain it. First, that attitude encourages a view that *everything is detail*, and that no idea is valid if an exception to it can be found. (In biology, exceptions can always be found.) Second, and more problematically, it imputes a bottom-up view of *what matters*, which is to say, of causation: we can't understand the causes of things until we have a comprehensive list of the parts involved. Those notions pervade, for example, current efforts to identify gene variants associated with traits by conducting statistical analyses of the genome sequences of hundreds or thousands of individuals. These studies do often find

correlations, and that's useful knowledge: it means we can make assessments of, say, the statistical risk a person has of developing a disease based on their genetic profile. But such studies are often accompanied with a warning that we don't *yet* understand how the gene variants cause the trait or risk in question. What they may overlook is that the correlation in question is that—a correlation—and nothing more. It might not be a sign of true causation at all.

As biochemist Mariano Bizzarri and his colleagues have said, “To understand mechanisms and provide conceptual insight into how and why processes occur, a shift of attention is required from genes to patterns and dynamics of the causal connections between components.” In this book I will try to talk about how life works in terms of causes. That, after all, is generally the sort of explanation science looks for, not least because it points to effective ways of intervening in a system to effect changes that can be predicted. In general, scientists won't be happy making claims about *how* some system works until they are confident that their explanation is predictive. In the book's final chapters I look at what might be done practically with the knowledge we have gained in recent decades about the way life—and especially the life that quickens our own bodies—functions.


### **Box 1.2: Cells as “Life's Atoms”**

Geneticist Paul Nurse, a 2001 Nobel laureate in physiology or medicine for his research on the cell cycle (the series of events that occur when cells divide), has said that the cell is “the basic unit of life.” There is no entity generally agreed to be alive that is not cellular.

And yet during the second half of the twentieth century the cell was destined to play second fiddle to the gene, which came instead to take pride of place in a reductionist dissection of life's processes. Only by degrees did it become evident that this will not do—for genes are not alive, but cells are. The cell, says Michel Morange, was “rediscovered” as “a major level in the integration of biological processes.” This rediscovery that we could not ignore cell-scale phenomena, he adds, “was probably the most important change that has taken place in

molecular biology since its rise in the 1950s”—and yet he says that most biologists didn’t even notice the change at all.

But not all cellular life is the same at this most fundamental level. As I mentioned earlier, the advent of modern microscopy revealed that cells have internal structure and organization. In our own cells there is a central compartment called the *nucleus*, separated from the rest by a membrane, in which most of the genetic material—the *chromosomes* containing almost all of our DNA—is sequestered. Our cells have a variety of other membrane-bound compartments, such as the *mitochondria* (where chemical energy is generated and where small strands of DNA contain thirty or so of our genes) and the *Golgi complex* (fig. 1.2).



**Fig. 1.2** A human cell is divided into many compartments, called organelles.

Bacteria, as well as another type of single-celled organism called *archaea*, have a different kind of internal organization. They have no nucleus; their DNA floats freely in the cellular liquid called the *cytoplasm*. Single-celled organisms that have no nucleus are called *prokaryotes*, and they are the most ancient form of life on Earth. Organisms whose cells have nuclei are called *eukaryotes*, and they are believed to have arisen later in evolutionary history—perhaps around 2

billion years ago—by the merging of a prokaryote with another cell, the likely nature of which is still disputed. There are single-celled eukaryotes, such as yeast, but all multicelled organisms are eukaryotes.

Multicellular animals, or *metazoa*, emerged during the Precambrian period; the earliest fossil evidence for them dates back to around 635 million years ago. As well as familiar bilateral animals like humans (which have mirror-image left and right sides), metazoans include sponges, jellyfish (*cnidarians*), and invertebrates called *ctenophores* or comb jellies. As we'll see, despite their shared cellular nature, metazoans have some distinctly different operating principles as compared to prokaryotes