CAMBRIDGE
Cambridge University Press & Assessment
978-1-108-83694-4 — The Biological Universe
Wallace Arthur
Excerpt
More Information

1 A TREE WITH MILLIONS OF TWIGS

The Tree of Life

The expression ‘tree of life’ is shorthand for four billion years of birth, death, reproduction, and relatedness. This extended family tree has been produced by four billion years of using energy from the environment to power biological systems. At present these systems, with which our planet is teeming, seem unique in the vastness of the cosmos. But they’re not. Their *apparent* uniqueness is an artefact produced by current limitations to human knowledge. One day we will have evidence of life on other planets, and that day may be close at hand. It’s not unreasonable to believe that our first evidence of extraterrestrial life will arrive in the next couple of decades.

In this book, our starting point for thinking about life in an interstellar context is the nature of life on Earth. Here on our home planet one particular tree of life has played out. This tree will continue to grow, though the directions in which its still-ungrown branches will extend are impossible to predict, so we cannot look with clarity into our evolutionary future. But we most certainly can examine our evolutionary past. And we can ask to what extent we would expect major features of that past to apply to trees of life that are playing out independently of ours – right now – on planets scattered across the Milky Way galaxy and beyond.

Notice that ‘tree of life’ is in the singular in the context of our own planet. Every living creature on Earth is related to every other. We humans are not just related to chimps, gorillas, and orang-utans. We are also related to the rest of the animal kingdom and, beyond that, to the trees we climbed as children, the yeast we use to make bread, and the bacteria that line our guts. The branches of the tree of life have no breaks in them. If we made a three-dimensional model tree of this kind, it would be possible to run a finger down from one terminal twig, such as humans, to a particular ancestor in the distant past, and then back up again to any other present-day twig, for example a maple tree.

But what shape should we choose when building our model? In other words, what shape characterizes the overall tree of life on Earth? It has been depicted in many ways since Darwin sketched an evolutionary tree diagram, in the form of lines gradually diverging from each other, in Chapter 4 of *The Origin of Species*. There are several caveats here, because the shape of the tree of life – or of parts of it – has been a source of heated argument among biologists over the years. So we need to tread carefully.

First, scale may be important. Let's consider this in terms of the two-dimensional trees that have been drawn on pieces of paper ever since Darwin. The shape of one small branch and the shape of the overall tree may not be the same. Second, at any scale we choose to examine, the divergence of branches may be leisurely (picture a V) or rapid (picture a U with a flat base). The former corresponds to a 'gradualist' view of evolution, the latter to either a 'punctuationist' or 'saltationist' view depending on the scale. Third, the vertical axis can represent time in an exact way, so that it could be labelled with units such as millions of years; or it could just represent time in a more general way in that it shows only the *order* of branching events, not their relative distances apart. Fourth, the horizontal axis could represent 'degree of difference' or it could be there simply to allow us to picture divergences – something that can't be done unless you have at least a two-dimensional diagram. The difference between these types of horizontal axis is that in the former case the distance apart of two twigs is a measure of their biological disparity, whereas in the latter it is not.

All the above four issues have been the focus of major debates at some stage in the history of evolutionary biology, and some of them continue to be debated. But the purpose of this book is not to examine such issues. We have a bigger picture to paint, so we'll sweep these issues under the proverbial carpet and focus on something even more important – the question of whether a tree diagram of *any* kind is the right way to depict evolutionary relatedness in the first place.

Consider for a moment an actual tree, whether a maple, an ash, or an oak. If you inspect it carefully in winter when no foliage obscures its branches and twigs, what you'll see are thousands of divergences but not a single convergence. Twigs grow apart from each other; they do not grow together and unite. But in the tree of life such growings-together do indeed happen to a degree. Two processes are responsible – interspecies hybridization and horizontal gene transfer. In the former process, two twigs, each representing a single species, hybridize and thus create a descendant species that is different from both of its parents. In theory this shouldn't happen, because a species is defined by its inability to interbreed with others – but in practice it does happen, because definitions are rarely perfect in the biological realm. In the latter process, DNA (deoxyribonucleic acid) from one twig is transferred into another, often via a virus.

The importance of these processes varies according to position in the tree. In the animal kingdom as a whole, their role is minor compared to twig divergence – though that does *not* mean that they aren't important. Some human genes appear to have originated by horizontal transfer from other species, including those as different from us as bacteria. Some of the best examples of interspecies hybridization come from the plant kingdom, while horizontal gene transfer is especially important in microbes.

How should we modify our picture of the tree of life to incorporate these two processes? Hybridization can be included simply by picturing twigs growing together – at least within some of the tree's branches. Horizontal transfer is probably better pictured as a sort of thin wire connecting two twigs at the same level (i.e. the same point in time). Taking both of these modifications on board (Figure 1.1), we now have a tree of life

Cambridge University Press & Assessment
978-1-108-83694-4 — The Biological Universe
Wallace Arthur
Excerpt
More Information

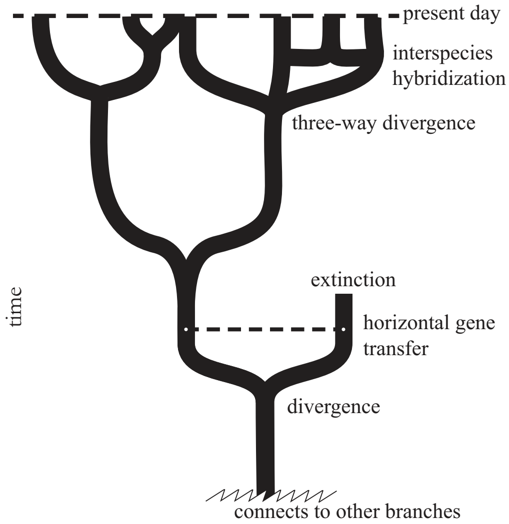

Figure 1.1 Part of the tree of life on Earth, showing divergences, interspecies hybridization, horizontal gene transfer, and extinction. All these features occur in any reasonably large branch of the tree, though their relative frequencies are expected to vary from one branch to another.

that is still largely tree-like but with some additional forms of growth compared to a real tree. A 2018 book by John Archibald – *The Tangled Tree* – provides further discussion of this issue.

The final thing to say about our tree-of-life picture (or model) is that its top should be flat. It's more like an African *Acacia* tree than a Norway spruce. This is because the present moment of time is the same for all the growing twigs, which collectively represent today's biota – the animals, plants, and other life-forms that populate the Earth right now. Let's take a look at this particular time-slice of the Earth's biological history.

Present-Day Twigs

So, now we alter our angle of view of the tree from the side to the top. We hover over it as a kestrel might, to achieve the proverbial bird's-eye view that we want. And we look at it as a

Cambridge University Press & Assessment
978-1-108-83694-4 — The Biological Universe
Wallace Arthur
Excerpt
More Information

photographer would when taking one of those shots where the foreground – in this case the present – is in sharp focus, and the background – in this case the past – is just a blur. We are then looking at a series of small circles, each one of them the tip of a growing twig. One is the human circle, another the bonobo circle, and so on. Species of cacti are represented by small circles far away from the ape ones. And mushrooms are represented by small circles far away from both of those other clusters.

Each circle is a species, though as we've already seen species can be badly behaved. The usual definition of a species is that while its members can breed among themselves none of them can breed with members of other species. And there is usually the proviso 'in the wild', so that we exclude information on what can happen in captivity, such as the production of ligers (lion–tiger hybrids). Of course, it would be naïve to expect all real organisms to conform to such a neat human concept. Some do, some don't. But even those that don't can be seen as fitting the definition in a probabilistic way – the density of reproductive interactions among members of a species is much higher than the density of such interactions between them and their sibling species.

Because there are at least a few million species on the Earth at present, and perhaps a few tens of millions, we need to have some way of structuring our knowledge of this vast biodiversity. And what better way than the method provided by the Swedish naturalist Carl Linnaeus in the mid-eighteenth century. Taking his approach, we group a bunch of neighbouring twigs together by drawing larger dotted circles around their small solid circles, thus representing groups of related species called genera (singular genus). For example, the orang-utan genus (*Pongo*) includes three twigs – those of the Bornean, Sumatran, and Tapanuli orangs (Figure 1.2). Our own genus (*Homo*) consists of only a single species in today's fauna. In contrast, some genera – for example the insect genus *Drosophila* – have hundreds of species.

In this exercise of looking down from above on the growing tips of the tree of life's twigs and drawing circles, we are doing something that can be described in terms of set theory. Our

Cambridge University Press & Assessment
978-1-108-83694-4 — The Biological Universe
Wallace Arthur
Excerpt
More Information

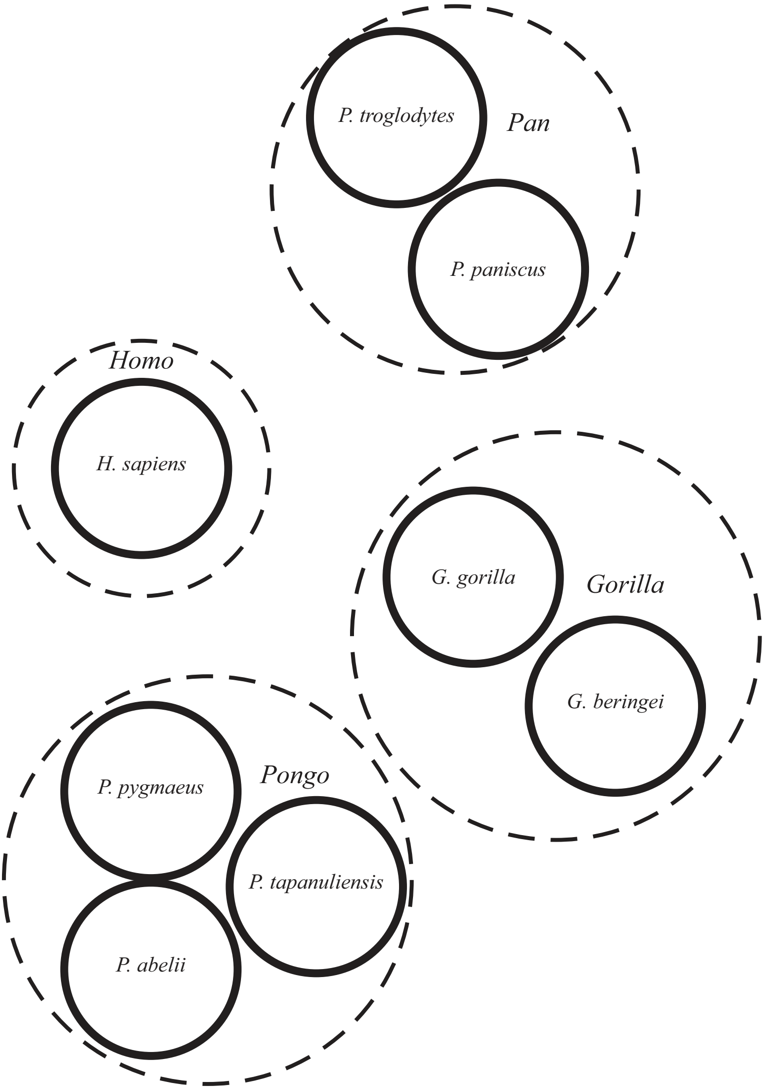

Figure 1.2 A particular part of the tree of life – the great ape branch – seen from above. Note the single extant species of humans (*Homo sapiens*), in contrast to the two or three species each of chimps, gorillas, and orangs. Common names are: robust or common chimpanzee (*Pan troglodytes*); pygmy chimpanzee or bonobo (*P. paniscus*); western gorilla (*Gorilla gorilla*); eastern gorilla, including mountain gorilla (*G. beringei*); Bornean orang (*Pongo pygmaeus*); Sumatran orang (*P. abelii*); Tapanuli orang (*P. tapanuliensis*).

circles-within-circles picture is what a mathematician would describe as a large set that includes one or more smaller sets. But there's something unique about our taxonomic sets: they are related to each other by their shared branches of the past. In a set of crockery types, where one subset is 'cups', another 'plates', and so on, there is no such underlying common ancestry – each item is made from scratch.

Taxonomists sometimes describe what they're doing as discovering and describing 'the pattern of natural classification'. The Linnaean approach draws bigger and bigger circles around progressively greater numbers of twigs, so that after species and genera we have families, orders, classes, and so on. Not only are these progressively more inclusive in terms of current biodiversity, but they are also progressively more deeply rooted in the tree of life. At the most inclusive end of the taxonomic hierarchy in Linnaeus's scheme was the kingdom – still in use today but expanded in number. Linnaeus described just two kingdoms of life – plants and animals. Now we also recognize at least one more – the fungi – and almost all biologists would say that there are several others. For example, all the large conspicuous brown seaweeds that we observe around our coasts, including those that make up that wonderful marine habitat called the kelp forest, are outside of the plant, fungal, and animal kingdoms. Studies on their genes make this conclusion clear. Collectively they are brown algae, but beware the term 'algae' as it has many inconsistent usages. They are in a fourth kingdom, even though there is some debate over what its name should be.

For Linnaeus, above kingdoms of life there was simply 'life'. But now we insert an even higher level of taxon than kingdoms – domains. The American microbiologist Carl Woese refined the taxonomic scheme of the earlier Carl in a major paper published in 1990. He grouped life-forms on Earth into the domains Bacteria, Archaea, and Eukarya. The first of these is self-explanatory, the last contains animals, plants, fungi, brown algae, and all other life-forms that are built of complex (eukaryotic) cells. The middle one, Archaea, was new, and based

Cambridge University Press & Assessment
978-1-108-83694-4 — The Biological Universe
Wallace Arthur
Excerpt
More Information

on earlier work by Woese and his colleagues. Superficially, the organisms that comprise Archaea look like bacteria, and pre-Woese they'd been classified as such. But, as he showed, they have a different form of RNA (ribonucleic acid) from the other two domains; and they use different fats in their cell membranes too. These are very deep-seated differences, and reflect their early divergence from both bacteria and eukaryotes.

It's important to realize that all the taxonomic categories above species – from genera up to domains – are arbitrary and have no clear definitions. They simply constitute a useful way of organizing information. The species is the only category that has biological meaning for the entities that comprise it, as opposed to for human observers; hence our ability to define, albeit imperfectly, what 'species' means. But there's an even more fundamental definition that we now need to consider – that of life itself.

What is Life?

If you'd like a lengthy discussion of this issue I can recommend the 2012 book of the same title by the organic chemist Addy Pross. Here we'll focus on just two approaches, which I'll call evolutionary and metabolic. The first is tightly linked to Darwinian natural selection. The second is linked instead to the biochemical processes that go on within cells. It's quite possible to be alive by one definition but not by the other; indeed, that's the case with viruses.

The evolutionary definition of life is as follows. Entities that exhibit the three properties of variation, reproduction, and inheritance are alive; those that don't are not. These are the very same three properties that are necessary for natural selection to occur. Consider a group of entities – we'll not prejudge the issue by calling them organisms just yet – that are rather similar but not identical to each other. They reproduce, in at least one of an immense variety of ways (beautifully discussed by Italian biologists Giuseppe Fusco and Alessandro Minelli in their 2019 book *The Biology of Reproduction*), and the offspring resemble their parent(s) more than they resemble randomly chosen members of the group. Resemblance in this context is not just external, nor just structural; it is internal and behavioural too. In such a situation, whichever variants are best suited to the current environmental conditions will leave most offspring, and so the composite nature of the population will change over time. Such a situation does not 'give rise to' natural selection – rather, it *is* natural selection.

At first this definition seems clear. According to it, birds and ferns are alive, while rocks and clouds are not. But if we dig deeper we find problems. Mules seem just as alive to me as do the horses and donkeys that were their parents. But, as sterile hybrids, they generally cannot themselves reproduce. Surely we shouldn't leap to the conclusion that they are inert entities, non-life-forms. And the converse problem of an entity that has the three requisite properties to be considered life but that we generally do not think of as life can also be encountered – for example computer viruses. These can exhibit variation, reproduction, and inheritance, but most of us would not consider them to be alive. And what about real (biological) viruses – are these alive? Many biologists see them as inhabiting a philosophical grey area between the living and the non-living. They can reproduce, but not on their own without hijacking another living system to help them. Then again, the same could be said of a tapeworm. So the evolutionary definition on its own is problematic.

The metabolic definition of life goes something like this. An entity is alive if it takes up energy and materials from its environment, uses these to maintain an internal state that is dynamic and yet buffered to some extent from environmental fluctuations, and ejects waste products from this process back into the environment from which the raw materials came. For the most part, this definition classifies entities as alive or not in the same way as does the evolutionary one: birds and ferns are alive, rocks and clouds aren't.

But again problems emerge when we start digging. The tiny invertebrate animals called tardigrades (or water bears) are famous for being able to withstand extreme conditions. They can survive extended periods of temperatures close to absolute zero, which would freeze-kill most other animals very quickly. They go into a state of suspended animation, from which they wake up when the ambient temperature is increased again. They use a similar technique to survive the vacuum of space; some of the tardigrades that have been taken into space – on the outside of a spacecraft rather than in the relative comfort of its interior – have survived and reanimated themselves on return to Earth, as reported by the Swedish scientist Ingemar Jönnson and his colleagues in 2008. Is a ‘cryptobiotic’ tardigrade alive? Personally I’d say yes, and that it’s just a rather extreme form of a hibernating hedgehog; but not everyone will agree with this view.

And what does being ‘buffered to some extent’ mean? We mammals can maintain an internal body temperature buffered into a narrow range around 37 degrees Celsius (98 Fahrenheit). Crocodiles can’t do that. Their internal temperature is much more variable over time – though it’s still buffered ‘a bit’ from the prevailing temperature of the environment. This more modest buffering is partly metabolic and partly behavioural in the sense of a crocodile’s choice of microhabitat.

What about the converse problem to that presented by the deep-sleeping tardigrade, in other words an entity that could be called ‘alive’ by the metabolic definition but which common sense would suggest is not? A fridge takes up energy from outside itself, uses this to maintain a regulated internal state, and ejects heat back into the environment as a sort of waste product. In this case we perhaps escape from definitional problems in that the fridge doesn’t take up *materials* from its environment. But even then a qualification is needed, because it does take up materials (cartons of milk, bottles of beer) – but not without human help, and it doesn’t use them to produce its internal homeostasis.

The metabolic definition could be modified by adding a stipulation that the ‘inside’ and ‘outside’ of the entity we’re looking at should be separated by one or more membranes – otherwise we conclude that the entity is inert. Again this works to