WALLACE ARTHUR

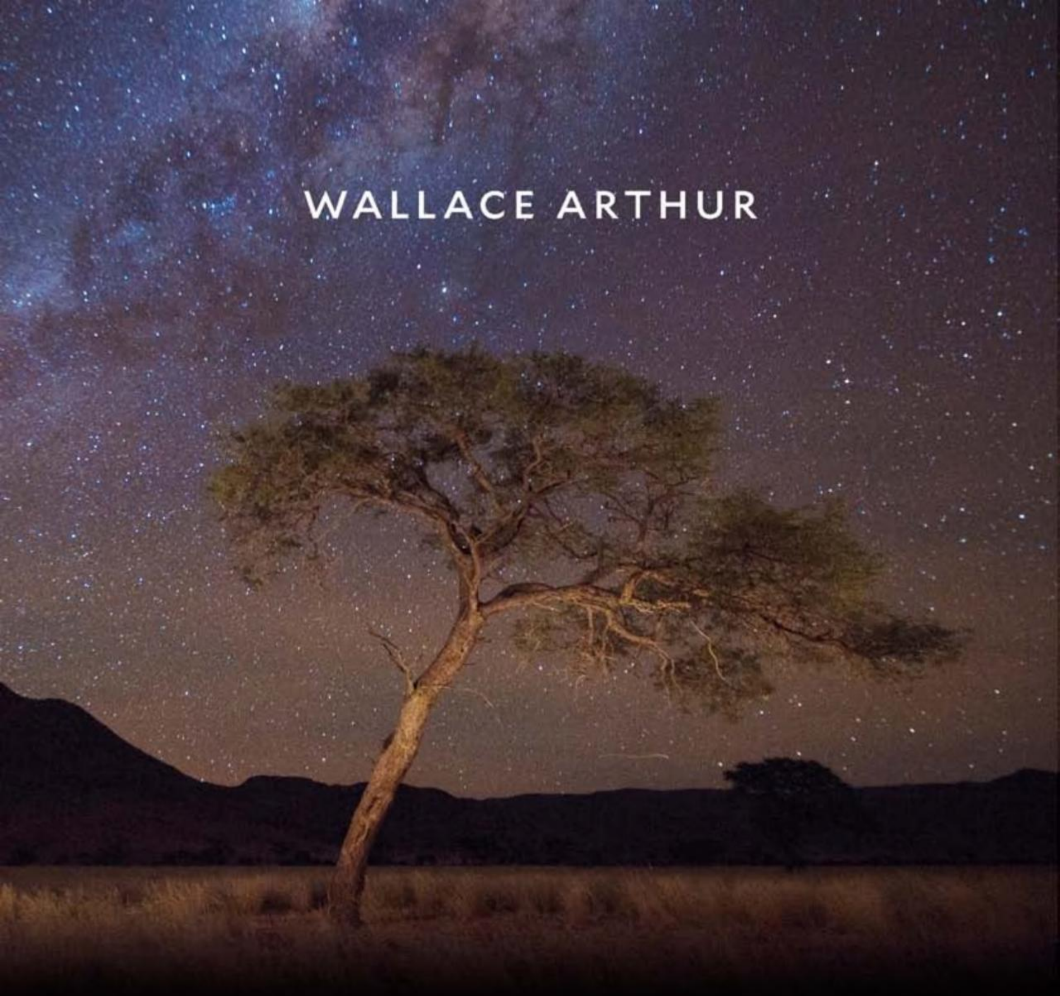

LIFE THROUGH TIME AND SPACE

LIFE *through* TIME *and* SPACE

WALLACE ARTHUR

Harvard University Press
Cambridge, Massachusetts & London, England
2017

Printed in the United States of America

First Printing

Library of Congress Cataloging-in-Publication Data
Names: Arthur, Wallace, author. | Arthur, Stephen, illustrator.
Title: Life through time and space / Wallace Arthur;
Illustrations by Stephen Arthur.
Description: Cambridge, Massachusetts : Harvard University Press, 2017. |
Includes bibliographical references and index.
Identifiers: LCCN 2016058797 | ISBN 9780674975866 (cloth)
Subjects: LCSH: Life—Origin. | Evolution. | Human evolution. | Life on other
planets. | Exobiology. | Extraterrestrial beings. | Developmental biology.
Classification: LCC QH325 .A724 2017 | DDC 576.8/3—dc23
LC record available at https://lccn.loc.gov/2016058797

Jacket design by Tim Jones
Jacket photographs © Minden Pictures: (top) camel-thorn tree in Namibia under
the Milky Way, (bottom) embryo of horseshoe crab

Contents

I. FROM STARS TO EMBRYOS

II. CYCLES OF LIFE

III. IN THE BEGINNING

IV. STRUCTURES AND FUNCTIONS

V. FROM BOULDERS TO BRAINS

VI. MILESTONES OF DISCOVERY

VII. ENDINGS AND ENLIGHTENMENT

Illustrations by Stephen Arthur

Preface

I've spent most of my life as a biologist; and in the last few years I've spent a lot of time studying astronomy. This book is a celebration of these sciences and of the growing relationship between them—a celebratory story told in plain language, not technical jargon. It's about our origins, our fates, our place in the universe, and the likelihood of intelligent alien life, looked at from both biological and astronomical points of view. I dedicate the book to all those who help to promote understanding rather than indoctrination, and in particular that great truth-seeker Thomas Henry Huxley. *Understanding* is one of the noblest goals of humanity, and one of the keys to our survival and progress.

The book is written as a series of seven triplets of chapters. Within each triplet, the first chapter is predominantly astronomical or astrobiological in flavor, the second is evolutionary, and the third is embryological. However, each chapter has arms that reach out into one or both of the other two domains. The connections between triplets, and between chapters within triplets, might at first seem cryptic, but they have a curious logic. I can best illustrate the nature of these connections by using the first triplet as an example.

In Chapter 1 we consider hypothetical (but maybe also real) inhabitants of the Andromeda galaxy, an object that can be seen in the nighttime sky by anyone with reasonably good eyesight. We use these extraterrestrial creatures as an entry point to the possibility of alien life in general. We imagine them looking toward the Earth with telescopes so advanced that they can actually see not just our planet but individual people wandering over its surface.

In Chapter 2 we meet some of the people they see. But these would not be us. Rather, since light from Andromeda takes 2.5 million years to reach us, and the same span of time to travel in the other direction, they would see protopeople of the distant past, belonging to an early species of *Homo,* characterized by a brain that's about half the size of our own.

In Chapter 3 we acknowledge that characterizing a species of human or protohuman as having a brain of a specific size is an oversimplification. As individuals, our brains are at first nonexistent, then small, then large. Early human embryos lack not only brains but any nerve cells at all. So, we contemplate the form of these early embryos, and the question of how they go about producing the beginnings of their nervous systems.

The other connections between chapters work in similar ways, so I won't waste words elucidating them further here. Their individual details will reveal themselves soon enough. The overall pattern of their linkage is designed to take you on a fascinating journey through embryological and evolutionary time, terrestrial and interstellar space.

Life through Time and Space

Chapter One

Galaxy Gazing

The Big W in the Sky

On a cloudless night, step outside and gaze up at the sky. What do you see? The short answer to this question is suns—or stars, which is a different name for the same thing. There are lots of them; just how many you see depends on where you live. If you're in a big city, you'll perhaps see only tens of them. If you're deep in the countryside, where the level of light pollution is low, you'll be able to see hundreds or, if you use binoculars, thousands. Not quite all of the bright objects you see in the sky will be stars. The exceptions will probably be a couple of planets, the lights of a few planes, and perhaps an orbiting satellite. What I want to direct your attention to, though, is a strange fuzzy blob which is none of these things. It's a galaxy, and one that almost certainly harbors intelligent life. When you look up at it, there is someone, or something, looking back at you.

The object I'm thinking of is the great galaxy of Andromeda. To find it, you can use a group of five bright stars that form the shape of a W. This is the constellation Cassiopeia, named after a vain queen in Greek mythology. It's one of the most conspicuous in the nighttime sky, if you live in the Northern Hemisphere. It can be seen year-round, even from quite light-polluted localities. Once you know that there's a big W in the sky, it's quite easy to find. Of course, a W is really two Vs stuck together. To locate the Andromeda galaxy, you use the right-hand V of this particular W as an arrowhead.

Here's what you do. Look at how deep the V is. Then project the angle of your view down about three times that depth, in the direction of the arrowhead and very slightly skewed to the right. There you will find the Andromeda galaxy.

How easy is this object to see with the naked eye? This depends on just three things. The first is your eyesight. If you have good sight, you should see it. If, like me, you have only average sight, you might see it and you might not—but with a pair of binoculars you'll be fine. The second is the level of light pollution. If you live in a large city, you might have to go out into the surrounding countryside to be able to see the galaxy. The third is the time of year. Although the big W is visible in all seasons, when you project the direction of your gaze the appropriate distance in the direction of the arrowhead, there is a time of year when this will take you to a point close to, or below, the horizon. For many of us, this will be just a short period in the spring; the exact duration depends on where you live.

Meet the Andromedans

Once you see the galaxy, consider the following. It looks like a little fuzzy blob, but it consists of about 500 billion suns (or stars). Orbiting most of those suns are planets, some of them much like our own Earth. On many of those planets, there are life-forms. This may at first seem like an overly strong statement: why "many" rather than "a few"? And why no "probably"? We'll soon see the answers to these questions; for the moment, please trust me that the sums work out in such a way that the likelihood of there being no life at all in the Andromeda galaxy is negligible.

Let's imagine two humanoid Andromedan scientists looking exactly in our direction, with a hypothetical telescope so powerful that they can see not just our planet but also individual humans walking across its surface. Even if such a device were possible, when these Andromedans look directly at us they do not see us. Why not?

Space and Time

To answer this question, we need to think a bit about space and time. The Andromeda galaxy is very close to us, as galaxies go. Admittedly it's many trillions of kilometers, or miles, away, but it's within our "local group" of galaxies (I love that phrase)—the ones that are really, really close to us, compared with all the others. Since neither kilometers nor miles work well for us when it comes to comprehending the vast distances of intergalactic space, we use other units, one of which is the light-year. We'll have to get thoroughly on top of this unit to understand exactly how far away Andromeda and other galaxies are from us. It has to become something that's as familiar to us as a meter or an inch.

The first thing to be clear about is that, despite its name, a light-year is a measure of distance, not time. You may already know this, in which case you'll probably also know that it's the distance light travels through space in a year. But how far is that? We can easily work it out. I was taught, as a child, that light travels at 186,000 miles per second. And so it does. But if you were taught using metric rather than U.S. customary or imperial units, then you'll have been told that light travels at 300,000 kilometers per second—which is the same thing, though a suspiciously round figure. If we know how far light travels in a second, it's easy to calculate how far it will travel in a longer period of time, like a year. To save you doing the sums, here's the answer: very roughly 10 trillion kilometers, or 6 trillion miles (these figures are not suspiciously neat; I've just rounded them to the nearest trillion).

Using this astro-friendly unit, how far away is the Andromeda galaxy? The answer: approximately 2.5 million light-years. Close enough to be "local" in astronomic terms (many galaxies are *billions* of light-years from us), but rather a long way from us in any other terms.

Now let's move from distances in space to distances in time. Actually, this is very straightforward, given that we're starting with the distance unit that we call a light-year. The time that light takes to reach us from Andromeda is, by definition, 2.5 million years. So when we look at the galaxy from Earth, we see it as it was 2.5 million years ago, when the light we're seeing right now was originally emitted from the galaxy and began traveling toward us.

Watching Our Ancestors

This looking back in time works both ways. So if those Andromedan scientists were looking in our precise direction last night, they won't have seen us. But they may have seen some protohumans, perhaps belonging to the species *Homo habilis* (literally, "handy man"; more information on these creatures will follow shortly).

I'm quite convinced that my hypothetical Andromedan scientists are real. Here's why. Observations made over the last couple of decades on stars / suns within our own galaxy—the Milky Way—show that many suns, not just our own, have planets. Not only that, but suns with multiple rather than single planets are common, and may well be the norm. Solar systems with one, two, three, four, five, six, seven, and eight planets are all known—but with the proviso that in each case the figure I quote is a minimum because other planets in the appropriate system may yet remain undiscovered. It seems likely that systems with more than our own eight planets exist too, and will be found soon.

Planets of Life

So there are lots of planets. But how many are Earth-like? The best guess at the time of writing is about 1 in 200, though this figure may well have changed a bit by the time you're reading this chapter, given that a typical book has a gestation period of about a year and the current rate of planet discovery is remarkably high.

This figure of 1 in 200 is reached as follows. We now know of about 4,000 confirmed exoplanets—the name given to planets orbiting suns other than our own. Of these, about 20 are Earth-like, though of course that leaves open the question of exactly *how* Earth-like. If this is a fair sample of our galaxy overall (the Milky Way is a bit smaller than the Andromeda galaxy), then, when we know more, the numbers 20 and 4,000 will simply be scaled up and the fraction of Earth-like planets will remain about the same. Assuming that the Milky Way and Andromeda are broadly similar in their composition, which seems likely, the fraction of planets that are Earth-like there will be approximately the same as it is here.

The route from the Andromeda galaxy to the likelihood of Andromedan life-forms works something like this. We'll guesstimate the number of planets in Andromeda as being the same as the number of stars—about 500 billion. That's probably an underestimate, but no matter; in fact, it's sensible to err on the cautious side when trying to estimate the likelihood of life. Now we can guesstimate the number of Earth-like planets as being 1/200 of this huge number, which works out to 2.5 billion. We'll be pessimistic about the fraction of these that embark on an evolutionary process that produces life—say, just 1 in 100, which is probably another underestimate.

This gives us 25 million planets with life. On what fraction of these has evolution produced *intelligent* life? Let's go with our 1 in 100 fraction again, so we're now down to 250,000 planets. So our issue: we remain somewhat in the dark about how galaxies were born. We'll come back to that particular type of origin later. For now, let's just say that if there was a common cloud from which we and the Andromedans came, it existed many billions of years ago.

Life on the Wing?

You may have noticed something interesting that snuck in untrumpeted in the previous paragraph. This was the idea of all the "constituent creatures" of the Milky Way—the implication being, of course, that there is alien life much closer to us than Andromeda. If that's true (it probably is, but we don't know for sure), why have I started out by asking you to consider the possibility of life so far away? The answer is that I want you to be able to look at one specific thing in the sky where we're pretty sure life exists. When that thing is a fuzzy blob that contains billions of suns, we can indeed be pretty sure. However, when it's a single star, the chances of there being life on one of its orbiting planets are actually quite low.

You'll recall that in the recipe for finding the Andromeda galaxy the first step was to locate the big W in the sky that we call Cassiopeia. The second was to consider the W as being made up of two Vs, the right-hand one of which we used as an arrowhead. The three stars of that arrowhead, from the right-hand edge inward, are Caph, Schedar (the spelling is somewhat variable), and Navi. Perhaps we didn't need to use these as an arrowhead to point to something else—perhaps these suns/stars may themselves have orbiting planets with life. So far, there is no evidence to suggest that they do. However, if we journey to another constellation in the northern sky, Cygnus the swan, there is a star called Kepler 186 that has a planet (186f) that is rather Earth-like and may well host life. This solar system is in the area of the swan's right wing. And there are many others in the same general direction.

The Excitement of Science

Now here's an important issue for what we call "popular science." The main aim of this difficult endeavor is to spread the findings, and indeed the *excitement,* of science, with a minimum of turgid detail. To *do* science, details are crucial. But to *learn* about science's big picture of things, they're not. Or, to be a bit more accurate, they can be minimized. And that's what I've been trying to do so far in this book, and will continue to do throughout, following in the tradition of others who have written in this genre.

But wait a minute: have I succeeded up to this point? Maybe not. Here's a list of the astronomical terms I've mentioned so far: *Cassiopeia, Andromeda, Milky Way, galaxy, light-year, exoplanet, Caph, Schedar, Navi, Cygnus, Kepler 186* (sun), *Kepler 186f* (planet).

That's already a dozen potentially new names. If you're an astronomer, probably none of them will actually be new. But for most people some will be new, and for some people most will be new. How can anyone commit to memory a list of new names without getting bored and losing sight of the big picture and the excitement of scientific discovery? It's vital to provide an answer to this question, for otherwise the mysteries of the universe will be eclipsed by detail, jargon, names. Any author guilty of achieving that appalling eclipse (and there are many) should be ashamed. I will try very hard not to fall into the jargon trap, though 12 potentially new names in fewer than that number of pages does not seem an auspicious start.

But there's a solution to this problem: replacing many individual names with a single framework on which to hang them. For the names we've encountered so far, here's such a framework.

Close, Middling, and Far

There are three domains of space: close, middling, and far. *Close* contains only our own solar system—the Sun, the Earth, the other familiar planets (Mercury, Venus, and so on), the Moon, and a motley collection of other things (asteroids, dwarf planets, comets). From the perspective of another star, such as Caph, this whole collection of stuff can be thought of as just a pinpoint in space. Indeed, that's exactly what it would look like from Caph—it would appear as a fairly ordinary "star" that, if magically zoomed in upon, would reveal all this extraordinary detail including, ultimately, humans.

Middling contains all the other stars of our Milky Way galaxy. Take the arrowhead of Caph, Schedar, and Navi, for example. Although the arrowhead and the W of which it is a part seem like flat entities in the night sky (as do constellations in general, because we can't really detect the third dimension of celestial depth), they're very far from flat indeed. Caph is about 50 light-years away. Schedar, at about 200 light-years, is roughly four times as distant. And Navi is approximately three times farther again, at about 600 light-years. However, in one important sense, these distances are all the same—that is to say, they're all middling.

To see how middling differs from close, consider this. The full span of our solar system (comets and other oddballs aside) from the Sun to the average orbital distance of the farthest-out planet, Neptune, is only a tiny fraction of a light-year—less than a thousandth, in fact. We don't even use the light-year as a unit of measurement at this spatial scale. But the closest star to us—in other words, the closest sun apart from ours—is more than four light-years away. So the closest star is more than 4,000 times as far away from us as is the farthest planet of our solar system. Truly, the close and the middling are different realms of space.

The same is true of the middling and the far. The nearest large galaxy to us, Andromeda, is more than 20 times farther away than the most distant star within the Milky Way. For the intergalactic distances of the *far* realm, we use either millions or billions of light-years.

The differences between the close, middling, and far realms of space can be illustrated with periods (or full stops for Irish, British, and other non-American readers). Here is our solar system and the nearest star:

.   .

The Sun and all the familiar planets, from Mercury out to Neptune, are well within the first period. If the nearest star to us has its own planetary system (we've recently discovered that it does), then all that stuff is within the second period. The nearest galaxy, thought of as an extension to this picture, would be a few kilometers / miles off the edge of the page.

Now, shrink our entire Milky Way galaxy so that it, with its billions of constituent stars, becomes a period. We can use the same mental picture to compare the distance between it and the Andromeda galaxy, as follows:

.   .

Note that this time the distance between periods is smaller, but it's still many times greater than the diameter of each. Also note that, with our galaxy collapsed, the close and middling realms are both now within the first period.

Another way of thinking about the difference between the middling and far realms of space is this. As you look up into the night sky and see individual stars of our own galaxy (like those of the arrowhead) and, close to them, a separate galaxy, consider the arrowhead stars and all the others of our own galaxy as raindrops on the windshield of a car in which you are driving along a country road toward a farmhouse light (representing Andromeda) that is just visible on the horizon.

Okay, enough of full stops and raindrops. The purpose of this exercise has been to provide a mental framework on which we can hang new names, in a sense putting them in their place and rendering them non-threatening. Thus, for example, it's enough to know that the arrowhead stars are in the realm of the middling; we can forget about their individual names and distances for most purposes. Likewise, it's enough to know that the Andromeda galaxy is in the realm of the far. That way, we can get an intuitive feel for that crucial third celestial dimension, the one we can't actually see. And we can appreciate that if we were on the surface of Mars, in the realm of the close, we would be so near to Earth that the use of our big W in the sky to find Andromeda would be almost exactly the same as from our normal vantage point on Earth.

Now let's have a closer look at what those Andromedan scientists saw when they used their amazingly advanced telescope yesterday to look at the Earth—*Homo habilis* and other early humans.

tohumans more than 3 million years ago, well before the earliest handy-person fossils.

Notice that I just referred to *Homo habilis* as “one lineage,” not “the lineage” of protohumans that existed about 2.5 million years ago. It’s hard for us to picture a world where there is more than one species of what might be called “people.” In the present-day world *Homo sapiens* is the lone representative of people; but in the world of *Homo habilis* there were at least two, and probably more. This raises the question of how we are related to them—and indeed how they were related to each other.

It seems clear that all the species of protohuman arose from a single humanizing lineage that split from the chimp lineage about 7 million years ago. So we’re all related to some degree. Exactly which species begat which other species is an issue that is still taxing the best minds in paleoanthropology. Here we’ll take the view that handy-person’s ancestry lies in the extinct species of southern apes called *Australopithecus afarensis*. Even if this turns out not to be true, what follows regarding the evolution of brain size is affected remarkably little.

From One Brain to Another

All the very early species of protohumans, including those referred to as southern apes, had brain sizes smaller than about 500 cubic centimeters (cc). This “marker volume” is the same size as the engine of a Fiat 500. I find it helpful to think of car engine sizes when dealing with brain sizes, as they provide the most common context for the use of cc to measure volumes for the non-scientist. Handy person had crossed this threshold, though not by very much. Its brain size is thought to have been in the range 500–800 cc.

Now let’s go forward rather than backward in time. More specifically, if we move to a mere half million years ago, we find a species called *Homo heidelbergensis,* named after the German city near which some of its fossils have been discovered. This species was probably also on the ancestral line to modern humans, though it's always important to stress the fact that our views on exactly who was ancestral to whom might yet change due to future fossil finds.

The brain size of this species was in the range 1,000–1,350 cc, the latter figure being also the average brain size of modern humans. Of course, we can now use liters instead of cc if we prefer. The 1.3-liter marker is interesting, as it is a common engine size in cars, but also a value within the ranges of brain size in both *H. heidelbergensis* and *H. sapiens*.

It should now be clear that although human evolution is complex and treelike, the same as the evolution of any other group of creatures, we've managed to retrieve a line, or lineage, by starting with handy person and focusing on just a few species, including one of its ancestors and two of its descendants. In a sense, we've mentally climbed our own evolutionary tree using a single route, or branch, corresponding to our special interest in the origin of *Homo sapiens*. We've ignored the other branches, but that doesn't mean they weren't there.

By the way, brain size is a very blunt instrument for measuring mental capabilities such as intelligence. It's a start, but only that. Also, we should remind ourselves that each species of protohuman had a range of brain sizes, not one specific size. And again, some individuals fall outside the range. The reason for this apparent contradiction is that the range of brain sizes quoted is usually one that applies to adults, whereas some of the fossils that have been discovered are juveniles. We'll come back to the developmental dimension of brain size soon.

Out of Africa

"Early" human evolution, a phrase I'm using for everything up to handy person, was different from its later counterpart not just in being below the 500 cc brain size marker but also in being restricted to Africa. However, after about 2 million years ago the evolution of protohumans began to be a much more global affair. Various species spread out of their African cradle, generally via the Middle East, and colonized much of Europe and Asia. However, of the species that did so, only our own *Homo sapiens* has survived thus far. Earlier African exoduses have been followed by extinctions. This applies to the well-known Neanderthals, and also to "Peking man," which belonged to different species of *Homo*.

So, geographically, Africa was our origin. This seems to have been established beyond reasonable doubt. But today we humans are found not just on the six habitable continents but also, thanks to technology, on Antarctica, whose land surface is not the permanent home of any species of mammal. We came from extremes of heat but have become able to cope also with extremes of cold.

The spread of our own species, *Homo sapiens,* was very recent when seen in the context of the whole of human evolution. Our ancestors are thought to have migrated out of Africa as recently as 100,000 years ago. When we're dealing with past eras that are measurable in thousands rather than millions of years ago, it's easier to feel the connection with the present. Smaller numbers of thousands project us forward: the first human cities were built in Mesopotamia (now Syria, Iraq, and Iran) about 10,000 years ago; almost 1,000 years ago, those Frenchmen that we call Normans, because many of their ancestors were Vikings or other Norsemen, invaded England. But we're all Africans at heart.

Feeling the Presence of Protopeople

One of the challenges I've set myself in this book is to try to collapse large distances in space or time so that we can *feel* the presence of creatures who live a long way away, or who lived a long time ago. In the case of any intelligent aliens alive today in our own Milky Way, in neighboring Andromeda, or in galaxies that are further afield, the trick is to try to mentally capture their simultaneity with us. As you read these lines, they are perhaps eating, sleeping, or walking to work. There may be a large void between us in space, but there's none at all in time.

To feel the presence of those protohumans whom the Andromedan scientists were watching last night, the trick is either the same or the opposite, depending on how we look at it. What we need to do now is to ignore the large gulf in time between ourselves and handy people, and concentrate on their closeness in space. So if we travel to the Olduvai Gorge in Tanzania, where many fossils of this species have been found (a journey of only a few hours in our time but much longer in theirs), and stand on a rocky outcrop that affords a panoramic view over the surrounding terrain, we are probably standing in the exact spot where at least one member of *Homo habilis* stood about 2.5 million years ago.

In each case, whether aliens or handy people, concentrating on the dimension—time or space—in which we are close to them is a useful technique to counteract the mind's natural tendency to concentrate on the other dimension, the one in which we are far apart. In both cases the use of this technique brings the distant creatures into full view so that we can see them in our mind's eye and picture what they may be doing—in one case right here and in the other right now.

There is, however, an important difference between imagining the aliens of today and the human ancestors of the past. The case for extraterrestrial life is based not on evidence (at least not yet) but on probabilities. If there is a one-in-a-billion chance that there is humanoid life on an unknown distant planet, and there are 100 billion such planets in a galaxy, then there should be about 100 different types of humanoid out there. In contrast, we actually have the fossilized bones and stone tools that belonged to our African ancestors.

Evidence is at the heart of the scientific endeavor, and without it we make no real progress. However, speculation is also at the heart of science, despite the disreputable (to many scientists) name of this mental activity. Indeed, speculation is logically prior to evidence. To see this, it's only necessary to consider the other, nicer names for speculation: wonder and (if it's rather specific speculation) hypotheses. Our natural curiosity leads us to wonder, and the wonder may then firm up into ideas about specific possibilities (hypotheses), which we then gather evidence to try to test. Don't let anyone persuade you that science is a simple search for evidence rather than the richer, many-sided endeavor it actually is.

What Were Our Ancestors Thinking?

We arrived at handy person by asking what those Andromedan scientists would have seen if they used their incredibly powerful telescope to look directly at us last night. Let's now engage in a little time travel (which I suspect is not possible) and sit among a group of *Homo habilis* in the Olduvai Gorge. What would we (and they) see if we had access to Andromedan telescope technology and could focus in on the Andromedan scientists' home planet? You know the answer by now: we would see their ancestors of 5 million years before the present. We have no idea what these would look like, though we can speculate that the course of their evolution might not have been so very different from ours.

As we prepare to time-travel back to the present, we might well ask ourselves this question: what would an individual handy man or handy woman have been thinking when he or she looked up at the stars? They wouldn't see even the ancestors of today's Andromedan scientists if they looked toward that galaxy, since they didn't possess any telescopes at all, let alone amazingly advanced ones. That galaxy would be a fuzzy blob to them, just as it is to us (though they'd need a variant on our arrowhead technique to find it because the shapes of constellations change over time periods of millions of years).

Chapter Three

A Human with No Nerves

Changes in Our Brains

Neither your brain nor mine stays the same for very long. Indeed, the essence of brain-ness is change. I suppose you could extend that statement to equating change with the essence of life in general: there's a huge difference between the rates of change of internal parts going on in any living cell and those taking place within an inert object such as a rock. But I'd like to distinguish between two types of change in our brains over time—one of which typically occurs in seconds, the other in years.

The *seconds* timescale applies to thinking. Right now, as you're reading this, those junctions between one brain cell and another that we call synapses are firing rapidly, enabling you to take in and process information. The *years* timescale applies to how your brain got to be the size it is now (about 100 billion cells) from a starting point of zero cells when you were very young—so young that your age would be negative if we used the normal point of reference, birth.

But let's use a different reference point—conception. This is the instant that *you* began, with the fusion of a sperm cell and an egg cell. From that moment, for about the first two weeks of your life, you had no brain, and indeed no nerve cells of any kind. Evolution took animals on a journey from no brains to big brains that lasted about 700 million years. The process of embryogenesis (and post-embryonic development) has powered our own individual journeys from the same start point to the same end point in about two decades. However awesome evolution may be, surely individual development is even more so?

Let's leave this philosophical question hanging and proceed in a more productive direction. It would be wise to consider cells in general before we consider nerve cells and brain cells in particular. There's nothing like having a broad base on which to build.

The Building Blocks of Our Bodies

Given that your body, mine, and those of all other humans are made up of tiny units called cells, these should be very familiar things to us. And yet they're not, at least in terms of firsthand viewing for a person who is not a professional biologist. Yes, we know they exist. And yes, cells have made their way out of the jargon-laden world of science and into everyday speech, for example in expressions like "he doesn't have two brain cells to rub together." But the problem is that cells are for the most part too small to see. The largest human cell is *probably* the quasi-spherical egg cell, yet even that is only barely visible—it's about a tenth of a millimeter in diameter. Nerve cells can be much longer than that but they're extremely thin. Hence the "probably": it's hard to compare the sizes of cells with very different shapes.

The adult human body is made up of many trillions of cells; there are more cells in your body than there are stars in the Milky Way. The brain accounts for about 100 billion of them. And there are further millions of nerve cells in the spinal cord and in the peripheral nervous system, which connects to every extremity of our bodies.

If what I've just said is true (which it is), who are these mysterious humans with no nerve cells at all? I've already given the answer in passing: very young ones. But let's dig into this rather brief answer and dissect it a bit. We need to *feel* the presence of these incipient humans, just as we *felt* the presence of aliens (Chapter 1) and our *Homo habilis* ancestors (Chapter 2).

After a few days of development, a human embryo consists of tens of cells—the stage that's referred to as a morula, which comes from the Latin word for mulberry, a fruit that looks a bit like a raspberry because it's a multiple fruit with lots of constituent spherules. If we think of these spherules as cells, we see the connection with early-stage embryos (see the illustration on page xii).

Counting Nerveless Humans

The mulberry stage lasts for about a week, roughly from day 4 to day 11 after fertilization. During this time it increases in cell number from about 16 to about 100. At no point during this early stage of our development are any of our cells nerve cells, though of course some can be thought of, later, as having been the *ancestors* of nerve cells. So at the very least the number of humans with no nerves at any point in time—for example, *now*—is the number of human mulberries. But how many is that?

We can get a rough estimate as follows. The total human population of the world is somewhere around 7 billion. We can divide it up into progressively smaller age classes. Starting with decade-based classes (e.g., all those people between 20 and 30 years old), there are about a billion people in each. Of those, about 100 million will be within a single year-class (say, age 20–21). In each month of that year (say, March birthdays) there will be about 10 million, and within one week of that month about 2 million. Taking a similar approach to unborn humans as we just did for twentysomethings, that's the approximate number of unborn humans with no nerves who are alive today, given that the mulberry stage lasts about a week.

But there are a few provisos. (If you aren't interested in them, skip to the next paragraph.) The distribution of people between age groups in any population of humans is not even. We can represent this distribution as a diagram in which the vertical axis is age and the span of the "block" for each age category (e.g., 20s, 30s, 40s) is the number of people alive in each. Usually the diagram is pyramid-shaped, with the highest numbers present at the youngest ages and a gradual diminution to the oldest. For example, there are more twentysomethings at any one time than there are octogenarians. Therefore, the number of human mulberries is higher than our rough estimate of 2 million. In addition, the stages before and immediately after the mulberry have no nerves either, which will distort our figure in the same direction. So probably 2 million should be regarded as a *minimum* figure.

Where are all these nerveless incipient people? Where is the nearest one to you right now? In a city center, she or he is probably within a radius defined by how far you could throw a smallish stone. Of course, due to the internal nature of human embryonic development, we never see them. With other types of animals that have external development, the hiddenness of early embryos is not an issue, but small size remains an obstacle to viewing. So we have very little direct evidence of the existence of nerve-free embryos of humans or other animals, but we know they exist. And, following from the ballpark calculation above, we have some idea of their huge numbers. We also know that each of us *was* a morula (or mulberry) many years ago, for there is no other track through time from fertilization to adulthood than the one that goes via this stage.

How Does Development Work?

So here's an interesting puzzle. From a starting point of a mulberry-like thing, wherein all the cells look pretty similar, how does this tiny brainless creature begin to make nerve cells? How on earth does it know what to do? Or, since it doesn't really *know* anything (you can't know things if you don't have a brain), how does the process of embryogenesis work? How does it almost always spurn the millions of wrong ways to make a human and choose the right way? Why does our brain always end up in our head and not in our left foot?

You may be thinking that the last paragraph consists of an overindulgence of questions. But if so, you'd be wrong. In fact, the number of questions the embryo has to ask—and answer—is vastly greater than the five I've just posed. However, for the purpose of studying development we can group different embryonic decisions together into those of the same type. This makes the job of trying to understand a very complicated process a bit easier. The mulberry has to do three main types of things, as follows.

First, it has to grow, which means that it has to multiply its cell number hugely. You cannot have a brain of 100 billion cells if you only have 50 cells in total. Second, it has to make different *types* of cell, for you cannot have a brain if you have no nerve cells. Third, it has to connect these two things. The multiplication of cells and their differentiation into various cell types have to be coordinated with each other so that, among other things, the brain does end up in the head (and the toes in the feet).

Here we'll focus on just the middle one of these three things—making different types of cells—and we'll look at just a single case of it, the making of the first nerve cells. This, after all, is the beginning of the route to our brains. Strange to think that so many of us know so little about it. Let's take a small step toward solving that problem.

Back to the mulberry, then. As it grows beyond about 100 cells, it gradually ceases to be a mulberry and turns into something else. Although the continuous process of development flows like a river, we try to understand it by inventing names for different stretches of the river, all the while knowing that one merges seamlessly into thin rod within our embryonic sausage secretes these blue balls. In the tissue immediately overlying the thin rod (picture this rod as being blue too, just to mentally distinguish it from the rest of the sausage) there is a high concentration of blue balls. Further away, the concentration subsides. Ectodermal cells receiving sonic hedgehog balls react not just to receiving them but to how many they receive, and this, among other things, determines what type of nerve cell they become.

You can picture this process going on indefinitely as a sort of relay race. Once a new type of cell has been formed as a consequence of "upstream" signaling, it in turn can go on to do some signaling of its own, and hence cause further "downstream" processes to occur in the embryo. Not all cell types do this, but some of them do, and that's enough to take development all the way through to its completion. It's complex, but in the majority of cases it works. And when mistakes are made, they're usually minor. Polydactyly (extra fingers), rare as it may be, is commoner than "polybrainy," which is either vanishingly rare or fictional, depending on your point of view.

So as we go forward through embryonic time, signaling between cells can be seen as an important process that drives development. But going backward, we encounter a problem. At the mulberry stage, all the cells seem more or less the same. So how can any one of them start making a signal while the others just end up receiving it? In other words, how does the process of cell differentiation *begin,* especially when all the cells contain exactly the same genes? It looks like a hopeless case.

But it's not. In fact, there are many ways in which an initially homogeneous embryo can begin to regionalize and make different cell types. One was proposed by the famous computer scientist Alan Turing, who also made important contributions to biology. Turing proposed a model wherein a ring of cells that were initially all the same could generate differences so that some of the cells would act as the starting point for tentacles, others not. Thus he could explain the origin of the tentacles that surround the mouth of the tiny freshwater creature that we call a hydra. And from the perspective of how developmental systems work, that's essentially the same problem as how we come to be sausage-shaped despite starting as a mulberry. The real reasons for the appearance of a hydra's tentacles or the shape-shifting that takes us from mulberry to sausage are more complex than Turing's model, but it was an important step in the right direction.

Beyond the Sausage

The dark strip down the dorsal side of our embryonic sausage that is our incipient central nervous system is a beginning rather than an end. This strip is initially flat or slightly curved, like the sausage-skin I'm using to portray it. It later turns up at the edges so that it becomes U-shaped. The ends of the U then grow together and fuse, forming a tube—the neural tube, as it's called in the embryo. Then, due to receipt of certain signals, the head end of the neural tube outgrows the middle section and the tail end, thus beginning the distinction between the brain and the spinal cord. And the brain cells just keep on multiplying until there are billions of them, so we can use them to write or read, as we're doing now, or for many other things. Mulberry becomes sausage becomes sentient life-form. An everyday miracle.

Now you can see that statements like "handy man had a brain size of about 600 cc" are very biased—in a way that the Italian biologist Alessandro Minelli calls "adultocentric," for example in his book *Perspectives in Animal Phylogeny and Evolution*. The brain of any human being is a trajectory through time. It starts at 0 cc and ends up much larger—exactly how large varies from one individual to another. This is as true of handy person or any of our other ancestors as it is of present-day humans. It's true of other animals too. We are all four-dimensional creatures; we have trajectories through time as well as space.

To us as observers, cells are very small things. But to a molecule of sonic hedgehog they're very big. So big, in fact, that each of these tiny blue balls in the cell is like the proverbial drop in the ocean. And yet to a chemist a sonic hedgehog molecule is *huge*. Compared to a molecule of carbon dioxide, many trillions of which you have exhaled in the last few seconds, it's an absolute monster. If a small Irish village with just a few houses around a crossroads is carbon dioxide, sonic hedgehog is New York.

Our bodies are made up of countless molecules, some large, some middle-sized, and some small. Each of these in turn is made of atoms that belong to an assortment of elements—for example, carbon, oxygen, and iron. Biologists trying to explain the near miracle of embryogenesis don't usually spend too much time thinking about where the elements come from. But astronomers do, because in the early universe none of them existed. So, where *did* they come from? Let's find out.

II

CYCLES OF LIFE