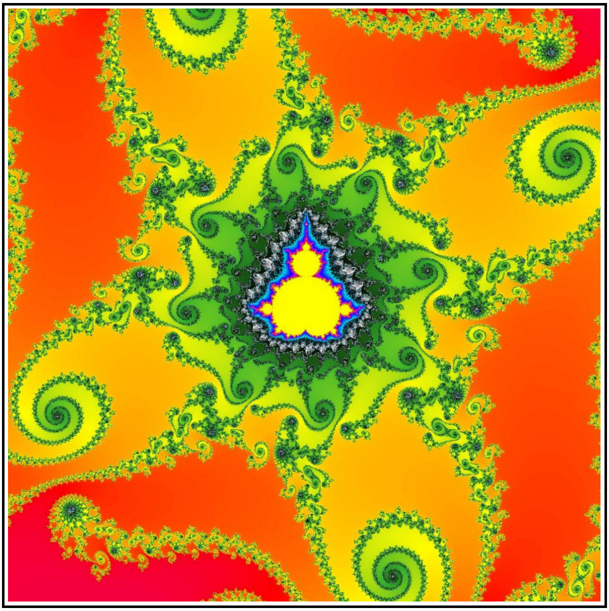

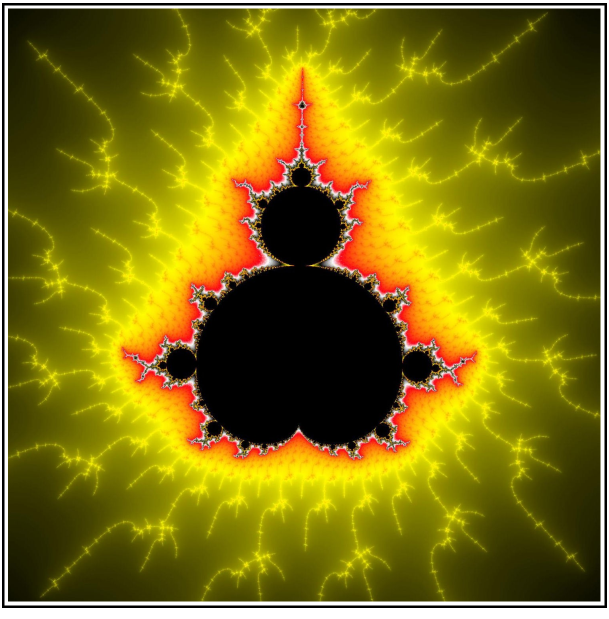

# The Mandelbrot Map

## a layman's guide

---------

Oliver Linton

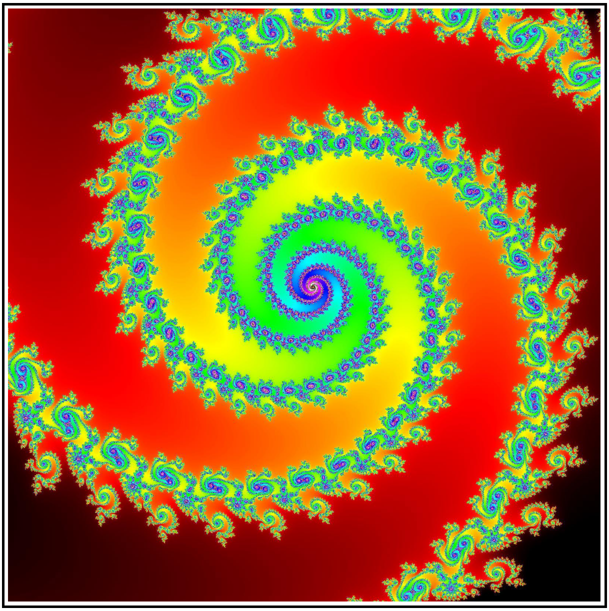

# Contents

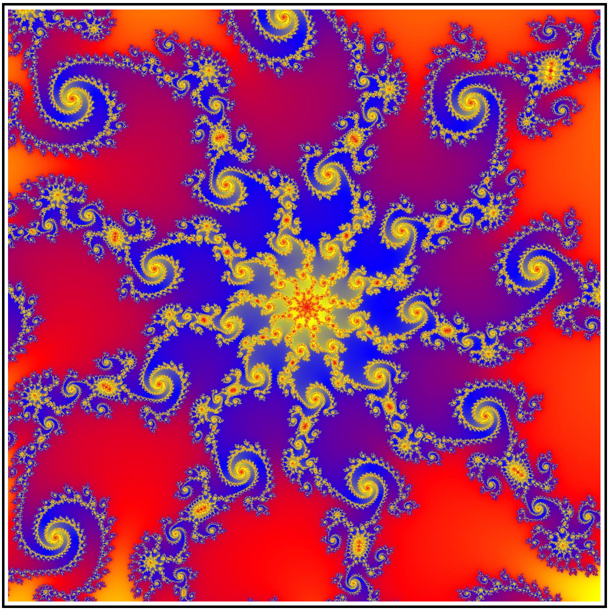

# Introduction

Many – one might even say all – processes in Nature are highly repetitive. Whether it is the orbits of the planets or the oscillations of a pendulum,  the patterns in the weather or the changes in the population of locusts, Scientists have always assumed that, fundamentally, these things are governed by more or less simple mathematical equations which are applied over and over again. This strategy has been remarkably successful in the case of the orbits of the planets and the oscillations of a pendulum which can be predicted with great accuracy; but predicting the weather or a plague of locusts has not proved to be so easy.

For many years it was thought that this was because the equations we were using were not sophisticated enough to model such complex behaviour but then in the early 70's when computers became powerful enough to make the dream of modelling the atmosphere in detail or investigating how populations develop a possibility, a remarkable discovery was made: it turns out that, under certain circumstances, even very simple equations can generate complex behaviour when they are iterated over and over again.

In 1976 the biologist Robert May published a paper in which he described the remarkable properties of the so-called Logistic equation:

$$x' = Ax(1 - x)$$

which can be used to model the annual development of a population of animals with restricted food resources. Normally such a population will find a stable value but May showed that if the birth rate was too high, the population would fluctuate wildly and apparently randomly. (For a detailed discussion of this remarkable equation see the companion volume to this 'Chaos and the Logistic Equation.)

A decade earlier the meteorologist and mathematician Edward

Lorentz had discovered similar behaviour in a set of differential equations which he was using to model the behaviour of the atmosphere. (More information about this and other 'strange attractors' can be found in another companion volume 'Fractals and how to draw them'.)

As long ago as 1887, the King of Sweden offered a prize to an mathematician who could answer the question of whether the solar system was or was not stable. Henri Poincaré won the prize but his answer was not very encouraging. Apparently, we cannot even be sure that at some point in the future the outer planets might not gang up on Earth and throw us out of the solar system (but the latest predictions tell us that this is unlikely to happen before the Sun blows up and destroys the Earth anyway).

At first, the study of such chaotic systems was regarded as a minor curiosity but when Benoit B. Mandelbrot published his book 'The Fractal Geometry of Nature' in 1982 it became an obsession with anyone who owned a personal computer with modest graphic facilities. One technique in particular proved to be absurdly simple to implement on a computer and produced stunningly beautiful results. The algorithm was called the 'Escape algorithm' and works like this: (For more details see: 'Fractals and how to draw them')

*Start with a complex number* $z_0$ *(the 'seed') and apply an iterative formula to it over and over again until the point either escapes to infinity or homes in on a stable cycle. If the point becomes stable, colour the point (x, y) on the screen corresponding to the complex number* $z_0$ *black, otherwise colour it in a shade which depends on the number of iterations needed for the point to escape beyond some arbitrary bailout value.*

The simplest and most widely studied formula is $z' = z^2 + c$ (where c is a complex number) and every value of $c$ generates a different image. These images are popularly known as 'Julia sets'. (Strictly speaking the Julia set is the set of all points which neither escape to infinity nor home in on a stable value but we will adopt the popular usage here.)

A couple of examples of Julia sets are shown below.

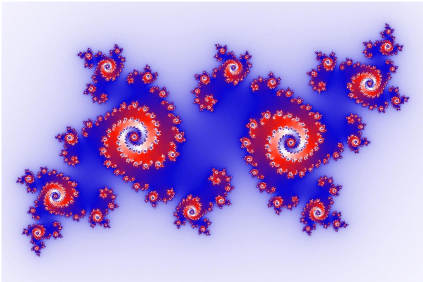

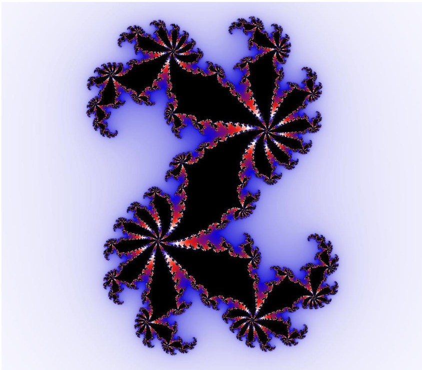

Some Julia sets, like the one at the top, are made of separate pieces and the point in the middle (which corresponds to a seed of (0, 0)) escapes to infinity more or less quickly. In contrast, other sets like the one below are connected and the origin is stable. In the latter case the the set is said to be 'filled'.

Using a suitable program it is possible to zoom in on different parts of the set but you will quickly realize that the image is self-similar. To put this another way – once you have zoomed beyond the point at which the whole image is visible, further zooming does not change anything.

Look at the top image and mentally separate the two halves. Each half is a distorted but nonetheless exact copy of the whole image.

We have already noted that for every value of ***c*** there is a unique Julia set. Here is a map of some of the Julia sets clustered round the origin

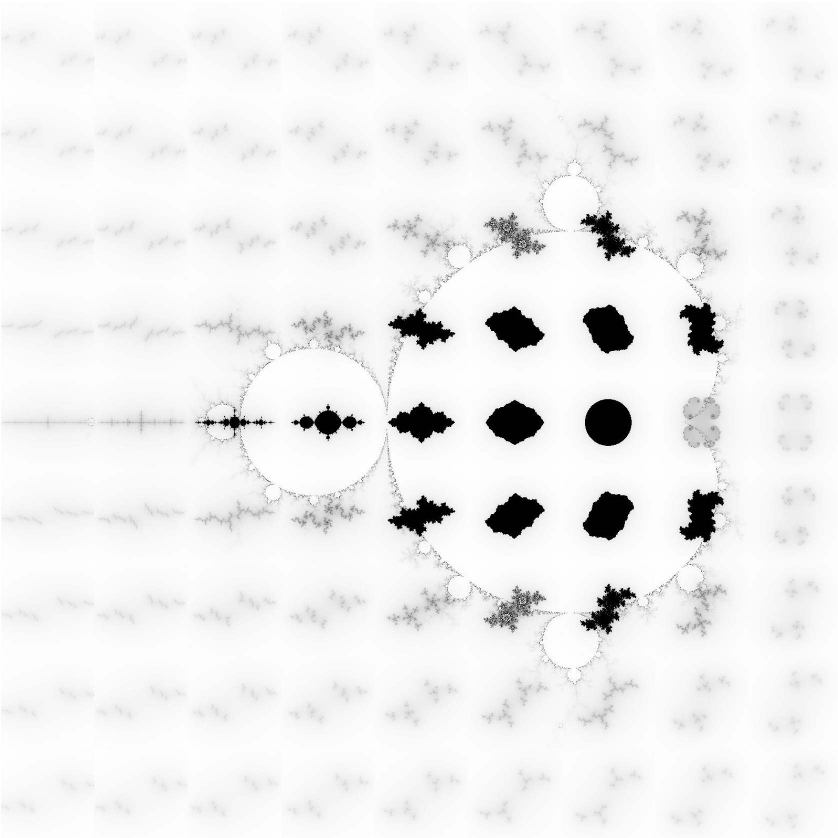

In the background you can see the shape of the famous Mandelbrot set. The Mandelbrot Set is a kind of map of all Julia sets. It is generated by iterating a fixed seed ((0, 0) in the classic case) over all values of ***c***. What this means is that the colour of the point on the Mandelbrot map is the same as the colour of the origin of the Julia set which corresponds to that value of ***c***.

# THE MANDELBROT MAP

The Mandelbrot Map is possibly the most intricate mathematical object known. It is so complex that zooming in to quite modest levels will get you to places that no one has ever seen before. It is also very easy to get lost. What we need is some way of identifying the different regions of the map. Fortunately, the map is highly organised and with a bit of practice it is possible to look at a certain small region and have a pretty good idea of where it is on the map.

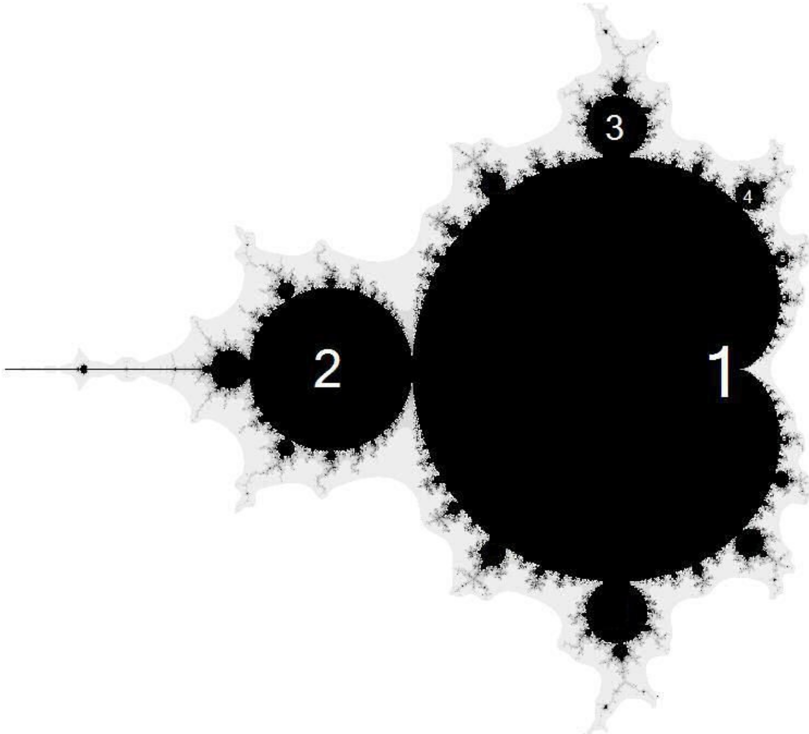


*The Mandelbrot Map with the principal lobes labelled*

# THE MAJOR LOBES

The most obvious feature of the map is that it consists of a lot of approximately circular 'blobs' stuck on other larger 'blobs'. Since the word 'blob' is rather vulgar, I shall refer to all the 'blobs' as **lobes.** The biggest lobe (which actually looks more like a cardioid) is lobe number 1. The next largest lobe (the 'frontal' lobe) is lobe number 2. The next largest lobe at the top of lobe 1 is labelled number 3 and subsequent lobes going clockwise to the right are labelled 4, 5, 6 etc. (Since the map is symmetrical about the X axis, we shall just consider the upper lobes for the moment. We shall see how to label the lower lobes later.) We shall call these lobes the **major** lobes and the whole sequence the **primary sequence of** lobes of which lobe number 1 is the first..

Now if you look closely at the main 'sprout' on lobe number 3, you will see that it soon splits into two branches at a major 'junction' which therefore has a total of 3 branches.

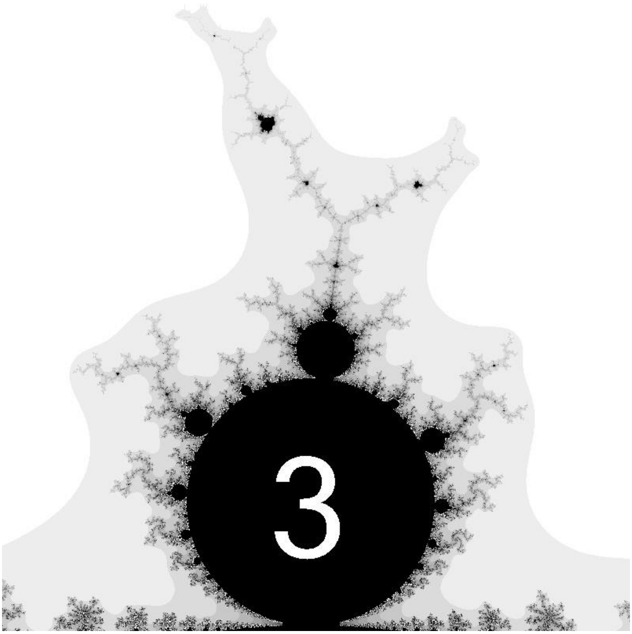

This is a good reason for calling this lobe number 3. I shall refer to the main 'sprout' which connects to the lobe as an **axon**; I shall call the main junction a **synapse,** the branches **dendrites** and the whole assembly a **neuron**. In this case the neuron has one axon and two dendrites and we shall refer to it as having **order** 3.

Now lets have a look at the upper right hand side of the main lobe in more detail:

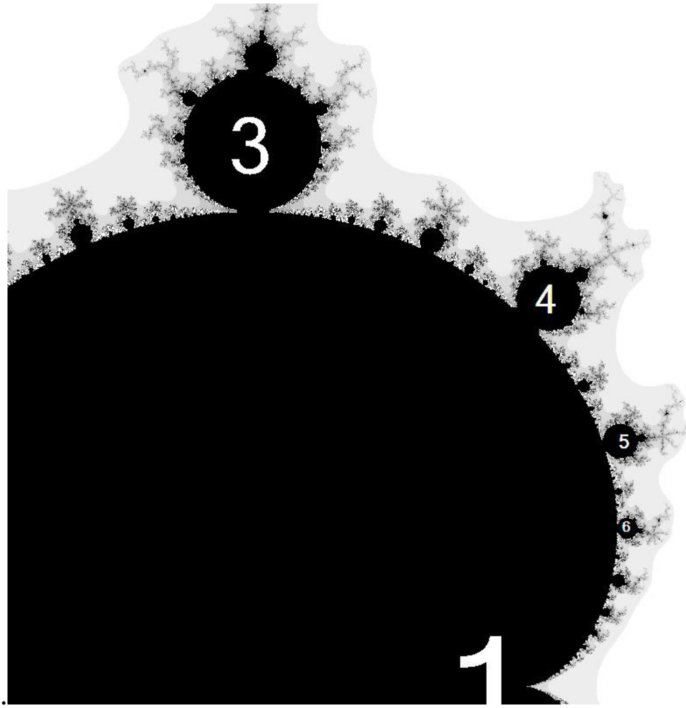


In each case you will see that the order of the principal synapse attached to each lobe is precisely equal to the number of the lobe.

The problem is – we have now used up all the integers – but there are still a huge number of lobes unlabelled. Let us see how these can be labelled consistently.

# THE MINOR LOBES

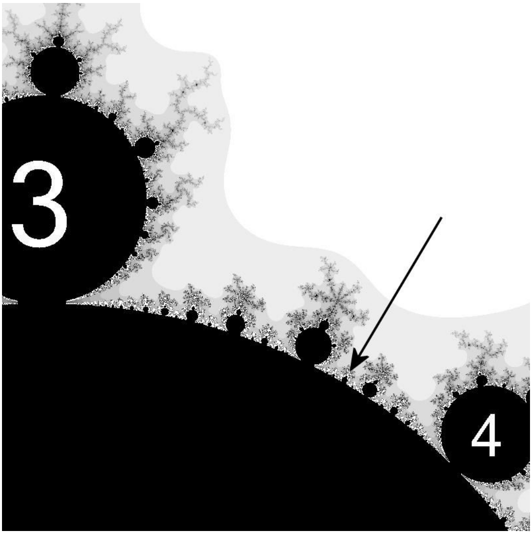


Approximately halfway between lobes 3 and 4 there is the next largest lobe. What shall we call it? You might be tempted to call it lobe 3›4 (where the angle quote means 'moving towards') being the next largest lobe between lobes 3 and 4 but there is a good reason why we should not do this. We have seen that lobes 3 and 4 are part of the *primary* sequence of lobes travelling clockwise round the main cardioid. Now the logical way to get to the lobe we are interested in is to go to lobe 4 first, then *change direction* to get to lobe 4›3. This is because lobe 4›3 is the second lobe in a *secondary* sequence of lobes starting at lobe 4 and going *anticlockwise* towards lobe 3. At first sight it might appear that between lobes 3 and 4 there are two sequences of equal status: one clockwise sequence starting at lobe 3 and another anticlockwise sequence starting at lobe 4, these two sequences crossing at the lobe in the middle. This is not the case. The anticlockwise sequence consisting of lobes 4, 4›3, 4›3›3, 4›3›3›3 … {3} etc. is a *secondary* sequence. (We can write this general sequence as 4›$3^n$) The clockwise sequence which we shall label 4›3, 4›3›4, 4›3›4›4 … {4} is a *tertiary* sequence and lobe 3 is definitely not part of it. (This sequence can be written 4›3›4$^{n}$)

I have marked one of the even smaller lobes with an arrow. What is its label? Well, to get there we should go first to lobe 4, then to 4›3, then another step towards lobe 4 will get us to 4›3›4, but what about the next step? You might be tempted to just add another 3 to the end but this is not correct for reasons which will soon become clear. The next step is taking us not towards lobe 3 but to lobe 4›3. The correct label is therefore 4›3›4›(4›3). Note that the brackets are essential. In fact it is good practice to include brackets whenever the flow changes direction e.g. ((4›3)›4)›(4›3).

Now lets turn our attention to the upper left hand side of the main lobe. The secondary sequence on this side stretches from lobe 3 to lobe 2. Its members are therefore labelled 3›2, 3›2›2, 3›2›2›2 ... {2} and its general form is 3›2$^{n}$.

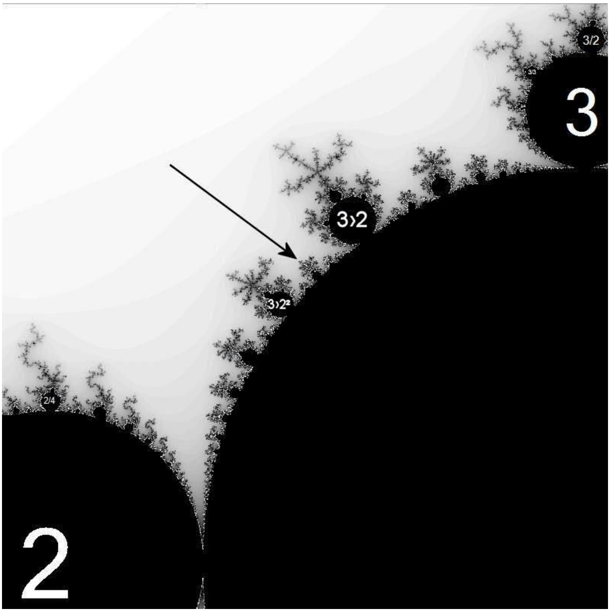


Can you work out the correct label for the arrowed lobe? (The answer is on the next page)

[1]Now if you use your Mandelbrot program to zoom in on the principal synapses of these lobes you will discover that the order of lobe 4›3 is 7 and the orders of the members of the sequence 4›$3^n$ are 4, 7, 10, 13 … (remember, the order of a lobe is the total number of branches round the principal synapse). It is immediately obvious that the order of a lobe is *the sum of all the numbers in its lobe label*. (e.g. the order of lobe 4›$3^4$ or 4›3›3›3›3 is 16.)

What about the order of ((4›3)›4)›(4›3)? If you zoom in on its principal synapse and count the branches carefully you will find that there are exactly 18 which is 4+3+4+4+3.

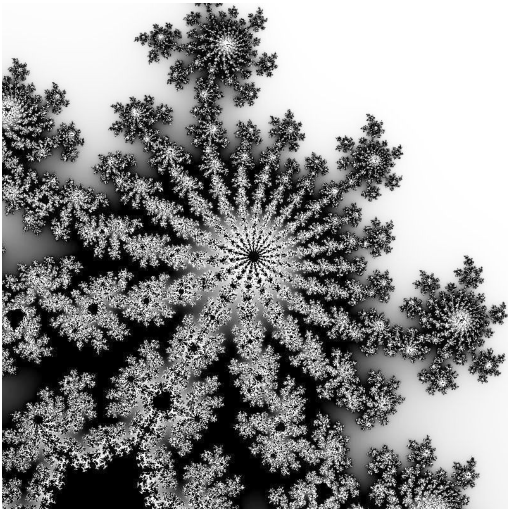

To summarise what we have discovered so far, every lobe can be given a unique label (the 'Linton label') which effectively describes a path of alternating clockwise and anti-clockwise sequences which have to be taken to reach the lobe. Every lobe is part of either a clockwise or an anticlockwise sequence; the clockwise sequences being primary, tertiary etc. and the anticlockwise sequences secondary, quaternary etc. and every lobe is the start of a sequence going in the opposite direction which terminates at (strictly just before) the previous member of the original sequence.

---

1 3›2›2›(3›2) or 3›$2^2$›(3›2)

We are now in a position to state the first Mandelbrot theorem:

- *The order of any lobe in a sequence is the sum of the orders of the previous lobe in the sequence plus the order of the lobe immediately beyond the limit of the sequence.*

This can be stated in a slightly different way as follows:

- *The $n^{th}$ lobe of a sequence which starts with a lobe of order A and terminates at a lobe with order B has order A + nB.*

This rule even applies to the primary sequence given that lobe 1 is both the start and the terminus of the sequence.

Now what about the lobes below the X axis? These constitute a secondary sequence which starts at lobe 2 and terminates at lobe 1. The lobes should therefore be labelled 2›1, 2›1›1, 2›1›1›1 etc. or in general 2›$1^n$.

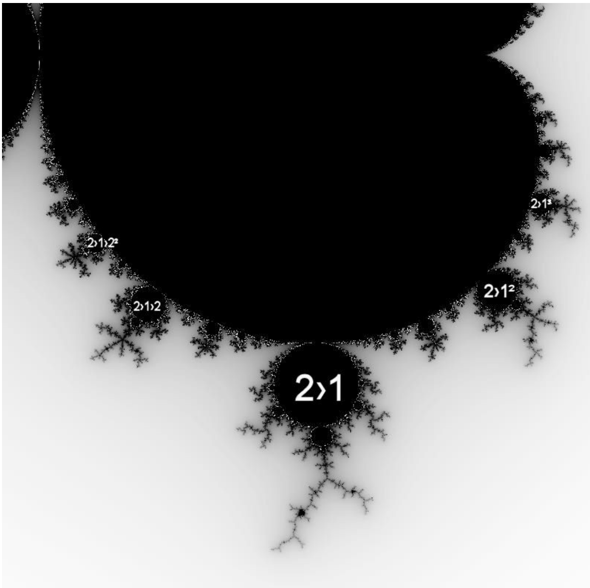


As you can see, the rule for calculating the order of the lobe works just as well for this sequence as it does for the primary sequence because 2+1 = 3; 2+1+1 =4; 2+1+1+1 = 5 etc.

# Order, Periodicity and Step Size

The Mandelbrot Map is generating by iterating the function $\boldsymbol{z}' = \boldsymbol{z}^2 + \boldsymbol{c}$ starting at the point $\boldsymbol{z} = 0$ for all $\boldsymbol{c}$ in the complex plane. When $\boldsymbol{c}$ is *inside* the main body of the map the point eventually settles down to a steady series of repeating values. When $\boldsymbol{c}$ is inside the main cardioid (lobe 1) it converges on a single point. When $\boldsymbol{c}$ is inside lobe 2 it jumps between 2 constant values. In general, when $\boldsymbol{c}$ is inside a lobe of periodicity (or order) $P$, the point jumps between $P$ values.

We have already seen that different lobes can have the same order – for example lobes 7, 4›3 and 3›$2^2$ all have order 7 and it interesting to trace out the path of $\boldsymbol{z}$ as shown in the following diagrams:

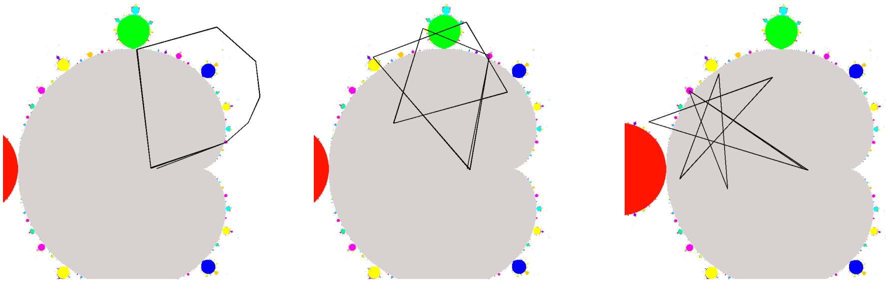

In the first case (lobe 7) the point $\boldsymbol{z}$ steps round the sequence one step at a time tracing out an irregular heptagon. In the second case (lobe 4›3) it steps 2 at a time while in the third it steps 3 at a time making a star shape. Now if you look at the star patterns generated by any sequence of lobes between lobes 3 and 4 you will always find that the step size increases by 1 every lobe down the sequence. Take, for example. The sequence 4›3›$4^n$. Lobe 4, being a primary lobe, has a step size of 1. Lobe 4›3 is one step down the sequence 4›$3^n$ so its step size is 2. Lobe 4›3›3 will have a step size of 3, as does lobe (4›3)›4; lobe (4›3)›4›4 has a step size of 4 etc.

So perhaps the step size is just an indication of how many steps you have to take to get to the lobe in question. Not so. Have a look at lobe ((4›3)›4)›(4›3). You will recall that this lobe has order 18. Now since it takes 4 steps to reach, you might think that its step size will be 4. In fact it is 5. The reason is as follows. When you step from 4 to 4›3 you are stepping towards a lobe (3) whose step is 1. So the step size increases by 1. Likewise when going from 4›3 to (4›3)›4 you are stepping towards another lobe of step size 1 so the step size increases by 1 again. But when you take the final step from (4›3)›4 to (4›3)›4›(4›3) you are stepping towards a lobe whose step size is 2 so the step size increments by 2. In fact, it is basically just the same as the rule for orders. This givers us the second Mandelbrot theorem:

- *The step size of any lobe in a sequence is the sum of the step sizes of the previous lobe in the sequence plus the step size of the lobe immediately beyond the limit of the sequence.*

  Or to put it another way

- *The $n^{th}$ lobe of a sequence which starts with a lobe of step size A and terminates at a lobe with step size B has step size A + nB.*

Now it may occur to you straight away that if the rule for step sizes is exactly the same as the rule for orders, then the step size of any lobe ought to be the same as its order. This is clearly not the case. But why?

The answer lies in the fact that lobe 1 is both the start and the terminus of the primary sequence of lobes 1, 2, 3, 4, … {1}. It is clear that the order of lobe 1 is equal to 1 – but what is its step size? Since an order 1 lobe doesn't step anywhere, we can assign the step size how we like. Let us assign it a step size of 1 for the purposes of starting the sequence off, but assign it a step size of 0 when it comes to adding extra steps in the sequence. Immediately all the members of the primary sequence will have a step size of 1. Everything else follows.

Calculating the step size from the Linton label is easy. Simply expand all the exponents and then count the number of numbers! e.g. the step size of lobe 3›$2^4$ will be [3›2›2›2›2] = 5. The step size of (4›3)›4›$(4›3)^2$ will be [4›3›4›(4›3)›(4›3)] = 7 because there are 7 numbers in the label.

When calculating the step size of lobes below the axis, for example lobe 2›1›1›1 which has a periodicity of 5, we find that it has a step size of 4. This makes sense because with 5 stepping stones, stepping round 4 steps at a time one way is the same as stepping round one at a time in the opposite direction.

So to summarise what we have said so far, every lobe has a unique 'Linton label' which tells us exactly where it lies on the main cardioid, and from the 'Linton label' we can calculate two numbers: the periodicity $P$ and the step size $Q$. The obvious question now arises: given the values of $P$ and $Q$, ($P > Q$) is there always a unique lobe which has just these properties. Amazingly, (provided that $P$ and $Q$ are co-prime) the answer is yes.

Let us see if we can find a lobe with periodicity $P = 8$ and step size $Q = 3$. We are looking for 3 numbers which add up to 8. In addition we should note that Linton labels only ever have two different consecutive digits in them (because any lobe must lie between two consecutive principal lobes) This means that the numbers we are seeking must be 2, 3 and 3. Now all we have to do is arrange these into a valid Linton label. The answer is, of course, (3›2)›3.

But what happens if we try to find a lobe with periodicity $P = 15$ and step size $Q = 6$ (noting that 15 and 6 have a common factor)? We start by looking for 6 numbers which add up to 15. These numbers are 2, 2, 2, 3, 3, and 3. But when we try to find a valid Linton label using these numbers we run into a problem. Obviously the label must start with 3›2. If we chose to change direction at this point we reach (3›2)›3 and we could go a step further to reach (3›2)›3›3. But now we reach the problem. We cannot go on because we have run out of 3's; but we cannot change direction either because that means moving towards (3›2)›3 making the next lobe ((3›2)›3›3)›(3›2)›3 not ((3›2)›3›3)›2›2. In fact there is no valid Linton label which uses 6 numbers adding up to 15. There is, however, a valid label which uses 2 numbers which add up to 5 namely 3›2.

So what is the deep reason behind all this amazing degree of organisation?

# ROTATION NUMBERS

We can state what we have learned in the form of a theorem as follows:

- *Every lobe is associated with a unique fraction* Q/P *where P is the periodicity of the lobe and Q is the step size and where P and Q are co-prime.*
- *Conversely, for every rational fraction Q/P expressed in its lowest terms there exists a unique lobe with periodicity P and step size Q.*

Now quite apart from the fact that the Mandelbrot Map is organised in the way I have described, it seems little short of miraculous that the two theorems which I have called the first and second Mandelbrot theorems should lead to the above result. The reason is as follows.

Every sequence of Mandelbrot lobes is characterized by a series of fractions called its *Rotation Number Q/P* where *P* is the periodicity of the lobe and *Q* is the step size. Lobe number 1 has periodicity 1 and step size 0 so the rotation numbers of the primary lobes are:

$$\frac{1}{1} \quad \frac{1}{2} \quad \frac{1}{3} \quad \frac{1}{4} \;\cdots \left\{\frac{0}{1}\right\}$$

The anticlockwise sequence 4, 4›3, 4›3›3, … {3} has the increasing series

$$\frac{1}{4} \quad \frac{2}{7} \quad \frac{3}{10} \quad \frac{4}{13} \;\cdots \left\{\frac{1}{3}\right\}$$

while the clockwise sequence 4›3, (4›3)›4, (4›3)›4›4 … {4} has the decreasing series

$$\frac{2}{7} \quad \frac{3}{11} \quad \frac{4}{15} \quad \frac{5}{19} \;\cdots \left\{\frac{1}{4}\right\}$$

In each case the next fraction is generated by adding the numerator of the limiting fraction (in curly brackets) to the numerator of the previous fraction (because of the second Mandelbrot theorem) and by adding the denominator of the limiting fraction to the denominator of the previous fraction (because of the first Mandelbrot theorem). (This method of combining fraction is called Farey addition after an 18$^{\text{th}}$ century mathematician called John Farey and the result is called the *mediant* of the two fractions).

Here is a map showing a selection of rotation numbers calculated from the Linton labels.

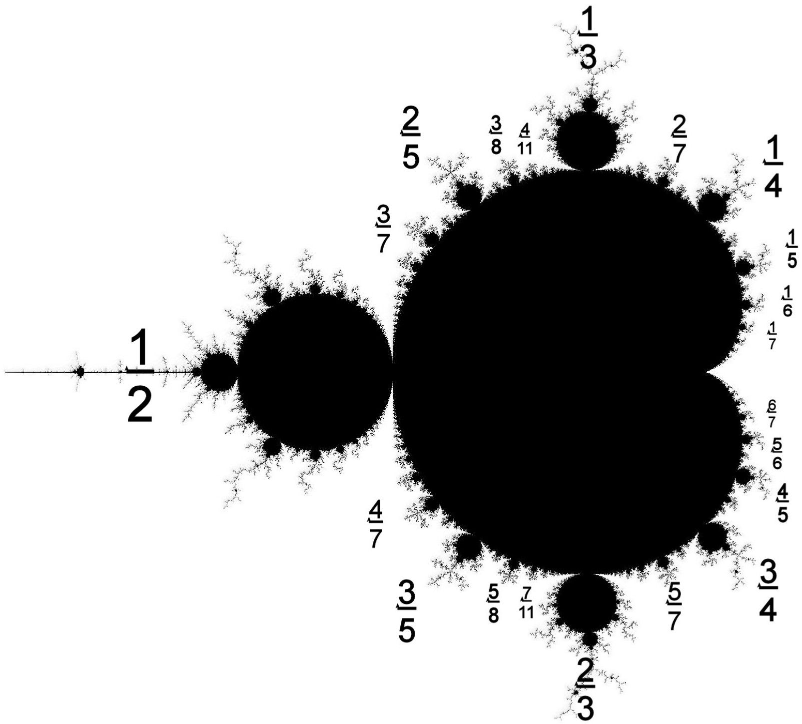


We can generate all these fractions by starting with the two fractions 1/1 and 0/1 and using Farey addition to generate a whole tree of fractions (called a Stern-Brocot tree) shown opposite. You will instantly recognise the rotation numbers of the lobes of the Mandelbrot Map (turned clockwise by 90°)

Now it is quite easy to prove that the Farey sum of two fractions must lie between the two original fractions and this guarantees that all

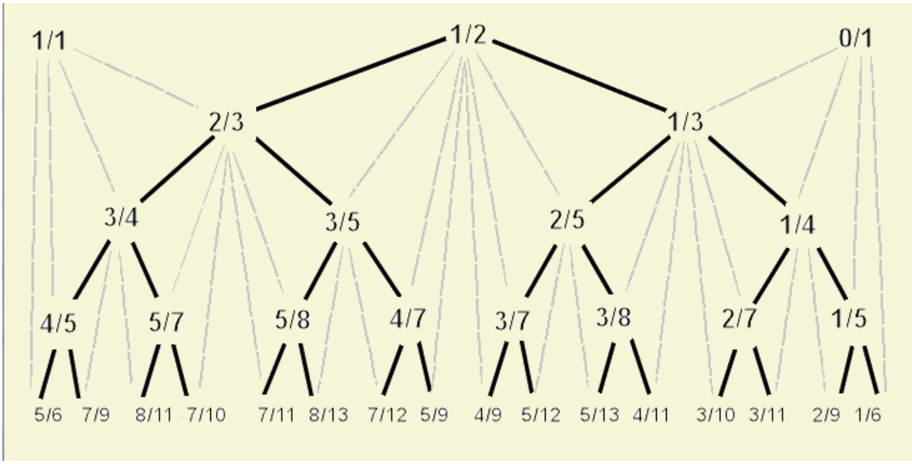


the fractions when read from left to right will be in numerical order. This fact alone proves that no fraction will occur more than once. This implies that every lobe round the main cardioid of the Mandelbrot map has a unique rotation number and that the rotation numbers are in strict numerical order starting at the cusp of the cardioid and moving round anti-clockwise.

What is not so obvious is the fact that *every possible fraction* (in its lowest terms) will be generated eventually. (The reason for this is basically that if there exists a fraction $a/b$ which does not occur in the Stern-Brocot tree, then you can calculate a second fraction $c/d$ which also cannot appear in the tree where $c<a$ and $d<b$. This means that you can calculate another fraction $e/f$ which cannot appear in the tree either and so on. Eventually you must reach the fraction $1/z$. But $1/z$ *does* appear in the tree so the assumption that $a/b$ does not appear in the tree must be false!)[2]

The Stern-Brocot tree gives us a way of calculating the Linton label from the periodicity and step size. For example, the lobe with rotation number 4/11 is reached from lobe 3 by taking one step towards lobe 2 and then 2 steps towards lobe 3. Its label is therefore (3›2)›3›3 and it is easy to confirm that its periodicity will be 11 and step size 4.

2 A slightly more rigorous proof is given in the appendix.

# Synapses, Periodicity and Step Size

What, exactly, are synapses? I have stated above that when $c$ is inside any of the main lobes (or any of the surrounding minibrots), the starting point $z = (0, 0)$ settles down to a periodic orbit whose periodicity is the order of the lobe (or minibrot). Now it is easily shown that for every value of $c$ in the plane, there are two values of $z$ which map onto itself (the roots of the equation $z^2 + c = z$). If $c$ is inside lobe 1, one of these 'self-mapping' points is an attractor and $z = 0$ homes in on it more or less quickly. (The other is a 'repellor' and is unstable.)

For any value of c, there are also values of $z$ which are periodic (e.g. the roots of the equation $(z^2 + c)^2 + c = z$ will have periodicity 2). If $c$ is inside lobe 2, the self-mapping point is unstable but one of the period 2 self mapping points is an attractor and $z$ oscillates between two points. We say that $z$ undergoes a *bifurcation* at the point where lobes 1 and 2 meet. Likewise, if $c$ is inside lobe 3, $z$ will home in on a cycle of period 3 etc.

If, on the other hand, $c$ is anywhere outside the principal lobes (or any of their sub-lobes), $z$ is very unlikely to find one of its self-mapping points and usually shoots off to infinity.

But just occasionally, after wandering around for a bit, $z$ may happen to hit precisely on one of its self-mapping (or periodically self-mapping) points. The values of $c$ for which this happens are called Misiurewicz points and these are the synapses. Like points inside the lobe they have infinite depth but unlike them they are repelling points not attractors. What this means is that if $c$ is even a tiny bit off the true value, $z$ rapidly spirals away from the stable point off to infinity. It is not surprising therefore that points in the vicinity of a synapse should have a degree of symmetry appropriate to the periodicity of the lobe to which it is attached.

I shall have more to say about Misiurewicz points and synapses when we come to examine the principal axon in more detail.

The periodicity of a lobe is obvious as soon as you look at the principal synapse but is there any way we can also determine the step size of the lobe and hence its fraction? Look at the two order 8 lobes shown below:

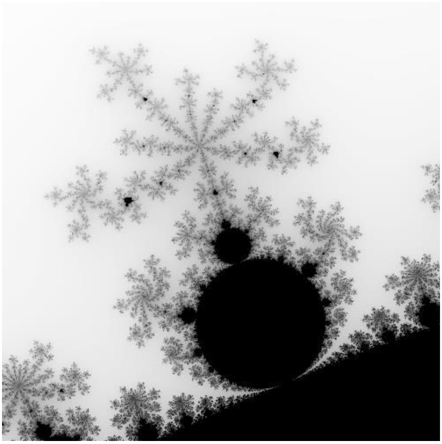

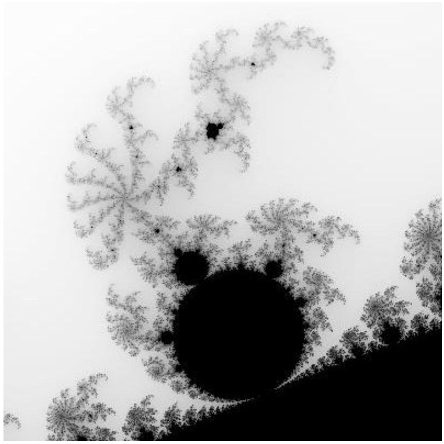

One of them is lobe 8 with associated fraction 1/8 while the other is lobe 3›2›3 with associated fraction 3/8. But which is which? The way to tell is by looking at the way the different dendrites are organised. In the image on the right, the dendrites increase in length steadily in a clockwise direction but in the one on the left the dendrites seem to be randomly organised. It is not difficult to guess that it is the one on the right which is attached to lobe 8. As for the image on the left, if you start from the main axon and step 3 dendrites at a time in a clockwise direction, the dendrites do generally get smaller and smaller but not consistently so.

There does not seem to be an easy way to determine the step size of a lobe unambiguously just by looking at its principal synapse.

# LABELLING TERTIARY LOBES

Lobe 1 is the *primary* lobe.

The *major* lobes 2, 3, 4 etc. are all *secondary* lobes because they sprout off the primary lobe.

Lobes 3›2, ((4›3)›4)›(4›3) etc. are *minor secondary* lobes – minor because they require several steps to reach but secondary because they still sprout off the primary lobe..

The question now arises: how shall we label the *tertiary* lobes which sprout off secondary lobes such as lobes 2, 3, 4 etc.?

It is obvious that all these lobes are organised in exactly the same way as the major and minor lobes which sprout off lobe 1. So all we have to do is use exactly the same notation as before but prefix it with an indication that the lobe referred to sprouts off lobe 2 not lobe 1. We do this by using a colon (meaning 'attached to') so the major lobes on lobe 2 are 2:2, 2:3, 2:4 etc.

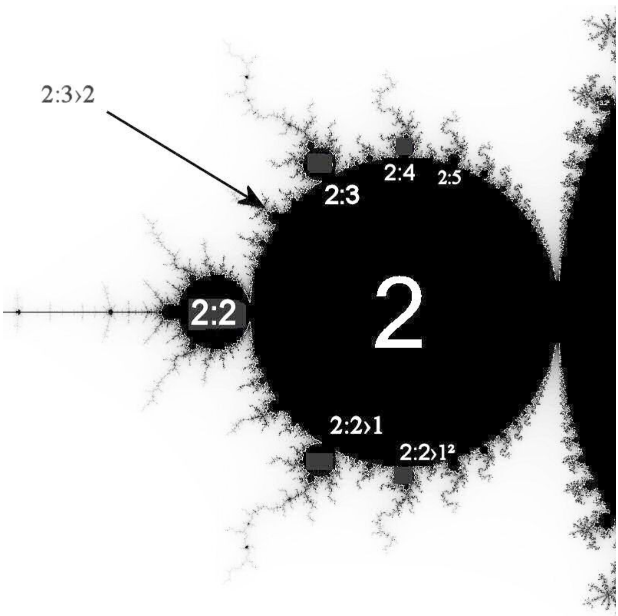

Now for consistency we should recognise that lobe 2 is not a primary lobe and therefore should strictly be called lobe 1:2. In fact all the labels we have used up to now should be prefixed with 1: – but we can often omit this without confusion. However, when referring to tertiary lobes, it would be as well to include it where appropriate.

For example, we might want to refer to lobe 3:2 (i.e. the largest lobe sprouting off the major lobe 3, illustrated below) and to distinguish it from lobe 2:3 (illustrated opposite). If we refer to these lobes as 1:3:2 and 1:2:3 respectively, it is more obvious which is which.

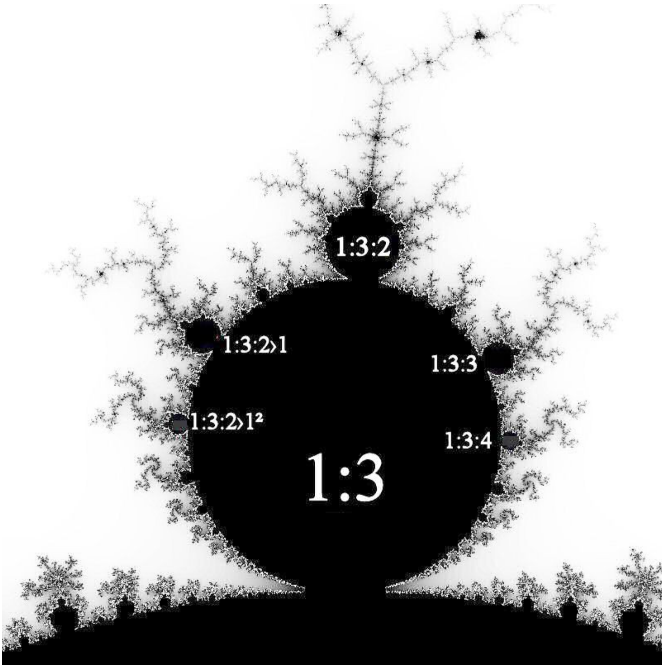


There is one final point which must be addressed. We have noted that the negative lobes on the primary lobe 1 should be labelled 2›1, 2›$1^2$ etc. and those on lobe 2 should be labelled 2:2›1, 2:2›$1^2$ etc. But what about the negative lobes on lobe 1:2:2? Should they be labelled 2:2:2›1, 2›$1^2$ etc. or 2:2:2›2, 2:2:2›$2^2$ etc. Are they heading for lobe 1 or lobe 2? The answer is clear when we consider the periodicity of these lobes.

# THE PERIODICITY OF TERTIARY LOBES

We noted that regarding the major lobes of lobe1, the order of each lobe (i.e. the number of branches in its principal synapse) was equal to its periodicity. This is not the case with the major lobes on lobe 2 because lobe 1:2:3 has period 6, in spite of the fact that its principal synapse has order 3. The rule, however, is clear. In order to calculate the periodicity of a lobe, you have to multiply the numbers on each side of the colon. In fact we can go further than this. If you have a lobe of order A sprouting off a lobe of order B the periodicity of the lobe will be A×B. For example, the period of lobe 1:2:4›3 will be 14 and the periodicity of lobe 1:3›2:4 will be 20. (Can you find this lobe?)

Now what about those negative lobes? According to the rules, lobe 1:2:2:2›1 should have periodicity 12 (2×2×3) while lobe 1:2:2:2›2 should have periodicity 16 (2×2×4). In fact it has periodicity 12 so the former appellation is correct. (Indeed the latter label does not really make any sense. How can you move *from* lobe 2 *towards* lobe 2?)

While we are looking at the tertiary lobes on lobe 2, it is worth noting that we can still assign the same rotation numbers to these lobes as the secondary lobes on lobe 1. In fact, not only do the rotation numbers determine the order in which the lobes appear round the lobe, they also determine their position. For example lobe 1:2:3 has rotation number 1/3 and is exactly one third of the way round the lobe. Likewise, lobe 1:2:4 is one quarter of the way round the lobe putting it exactly at the top. (You are probably wondering why this simple relation appears to hold for the tertiary lobes on lobe 2 but not the secondary lobes on lobe 1. Have patience!)

What can we say about the step size of these tertiary lobes? The following diagram shows the orbit of $z$ when $c$ is inside lobe 1:2:3. It has periodicity 6 but it does not step round the 6 points in a cyclic way; instead it steps alternately round two groups of 3 points. Its rotation number is still 1/3 but we have to bear in mind that we must interpret the numerator slightly differently.

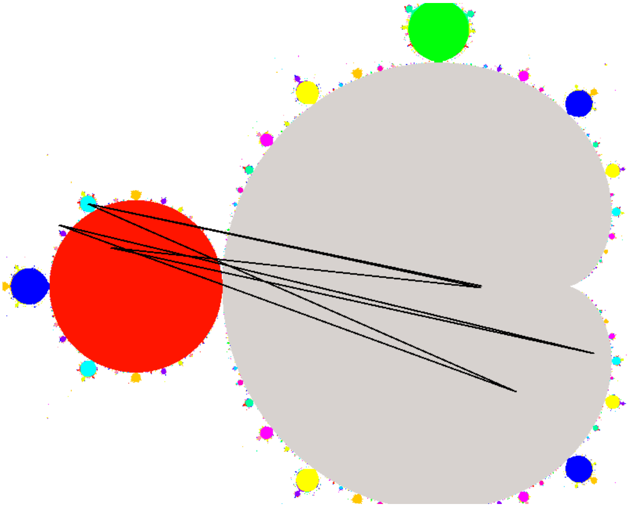

Would you care to hazard a guess as to what the orbit of $z$ would look like if $c$ was inside lobe 1:3:3? Well, surely it will step round in 3 groups of 3. Here it is:

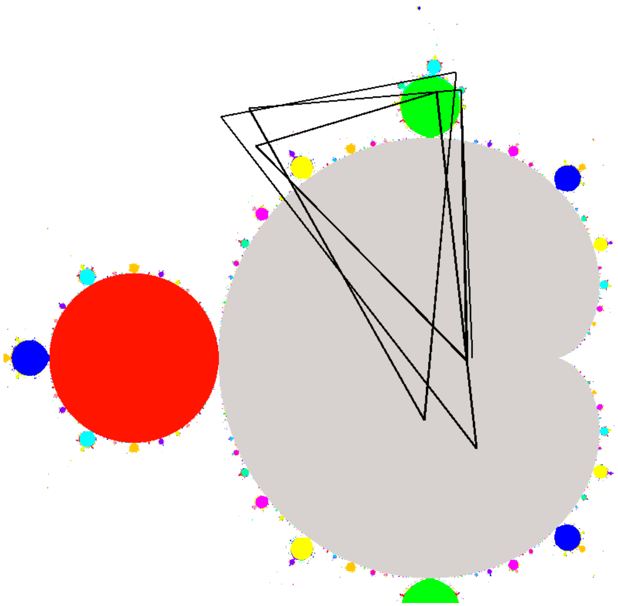

# Neurons and their Features

Now we know which parts of the map we are talking about, lets start to look at the various different neurons and their features. Here is the principal neuron of lobe 3:

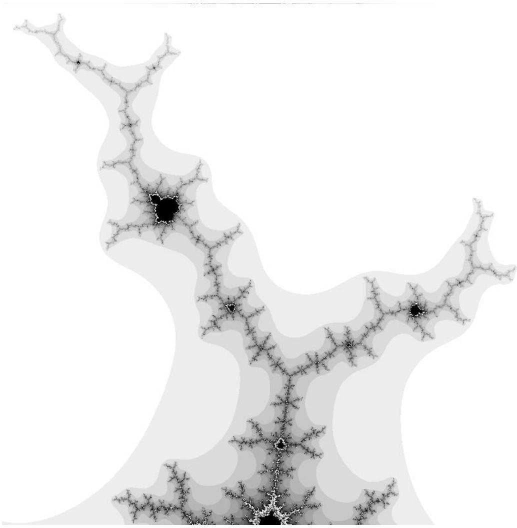

Of the two dendrites it is the first (measured clockwise from the axon) which is the longer and leads to the largest minibrot. (Minibrots are small islands of stability with the same general shape as the whole set.) The axon has a minibrot approximately half way between the lobe and the synapse and, if you look closely, you will find an infinite number of other smaller minibrots on either side making a fractal sequence. The whole structure is linked together with a filigree of channels which are believed to have infinite depth. Every single synapse has order 3. You can magnify any synapse as much as you like and you will never reveal anything other than a straightforward junction. If, however, you spot what appears to be a synapse with a higher order – for example, 6 – you will find a minibrot there with two order 3 synapses, one on each side. This is one of the smaller minibrots on the principal axon of lobe 3:

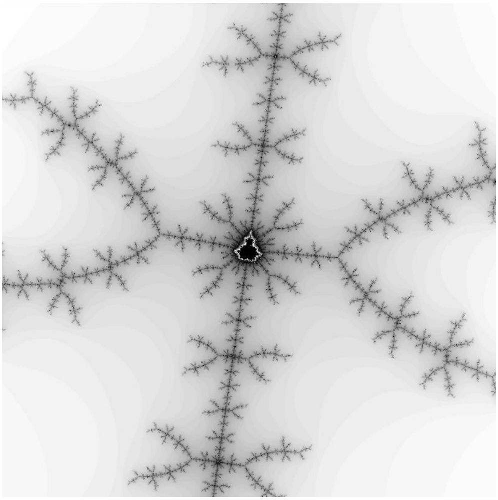

In fact, every minibrot in the map has this basic structure: an axon (or dendrite) running through the middle of the minibrot and two more or less prominent synapses symmetrically on either side. You will notice that four smaller order 3 synapses have appeared in the four quadrants – and eight even smaller ones between them. As you magnify smaller and smaller minibrots, the binary symmetry grows so that eventually you can find minibrots surrounded by 64 or 128 synapses like this one:

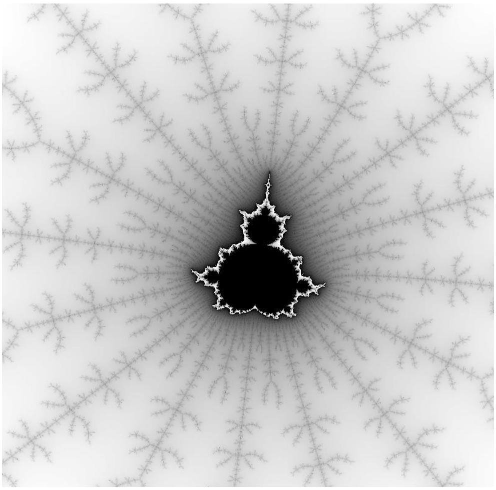

Neurons of high order lobes are much more interesting. This is the principal neuron of lobe 15 which is in the region known as 'elephant valley' for obvious reasons (though this elephant happens to be upside down!):

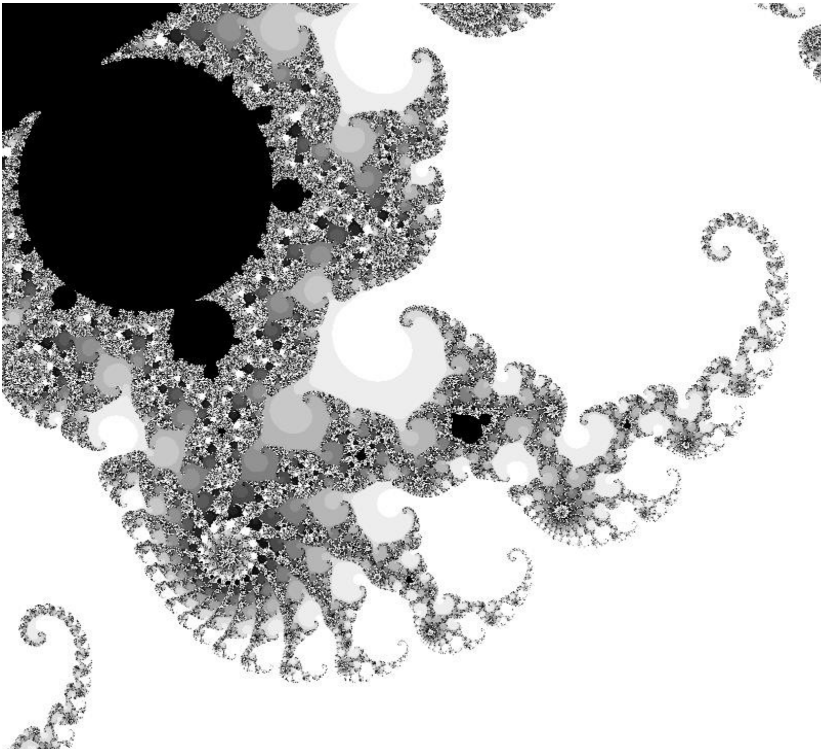

The principal synapse is the spiral structure at the bottom left. The first dendrite (measured clockwise from the axon) is the longest. Halfway along there is a large minibrot. If you trace along its principal axon (the one that emerges from lobe 2 on the minibrot) you will see that it leads to another synapse of order 15 but this time it is dendrite number 14 which is longest and leads via a small minibrot to the next largest synapse. This pattern is repeated continuously on a smaller and smaller scale resulting in a beautiful spiral which seems to disappear down into infinite depths. Every dendrite on every synapse ends in one of these **infinite spirals**.

The image below shows the principal neuron of lobe 3›2[6] which also has order 15 and step size 7. (You will notice that the 14 dendrites fall into two distinct series of 7. This is, of course, related to the step size of the lobe.) The largest dendrite ends in an infinite spiral forming a curving head which has reminded some people of a sea horse. This wedge-shaped promontory between lobes 1 and 2 is often known as 'sea-horse valley'.

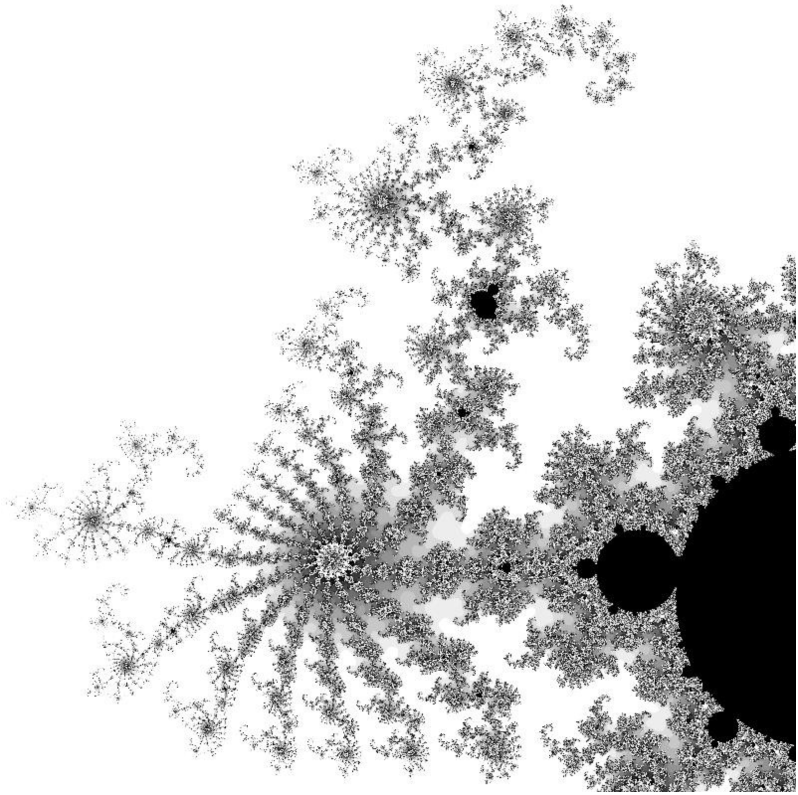

The minibrots which link the large synapses together are strikingly symmetric and would make a good design for a wallpaper. (Once again, it is worth noting the basic structure of a minibrot, namely an axon or dendrite running through it with two prominent synapses, one on each side.)

The images on this page and the next are of the principal neuron of lobe 1:2:12 which can be found on the opposite side of 'sea horse valley'.

You can tell that they are on lobe 2 rather than on the main lobe because the infinite spirals are double: they go in and then come out again! This gives the linking minibrots a particularly pleasing aspect. (The point at the bottom of the spiral is another of those Misiurewicz points and is, in fact, a synapse of order 2 – but I will say more about this later.)

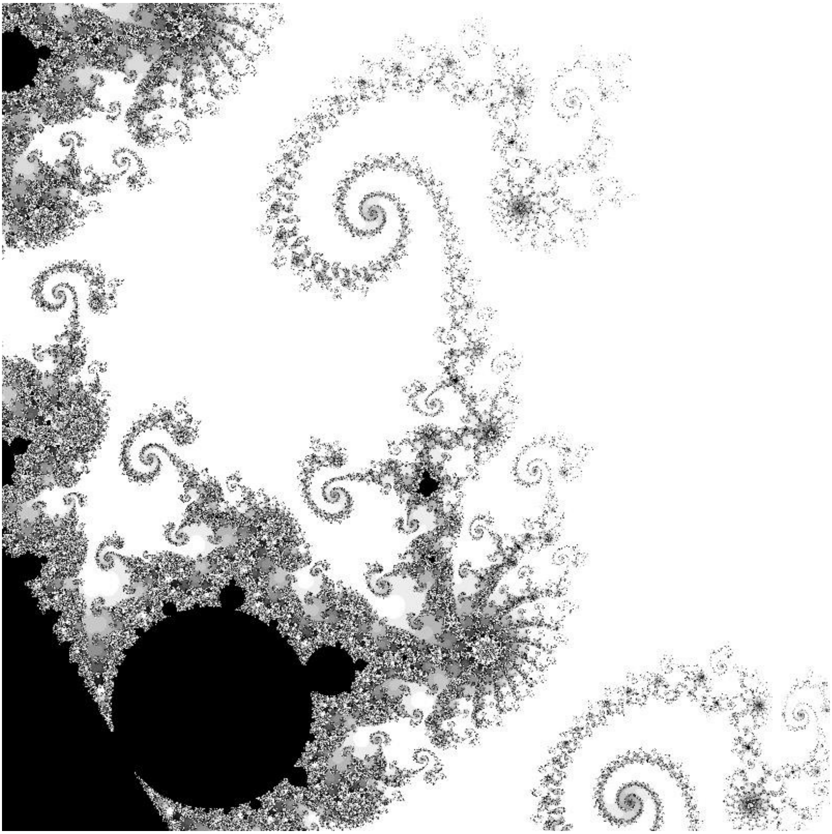

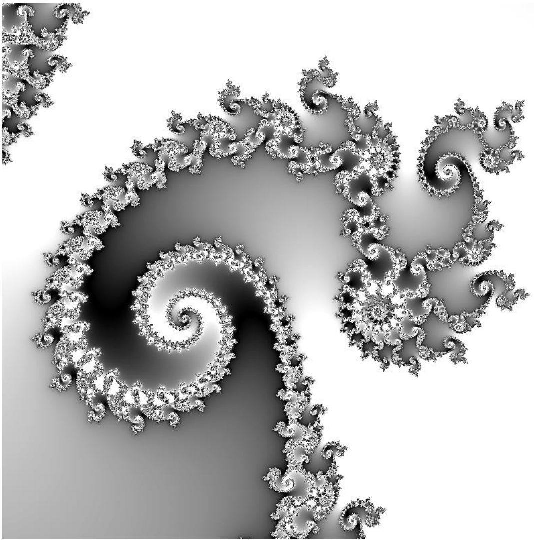

I love this image – it looks to me like a beautifully ornate bracket for a candlestick!

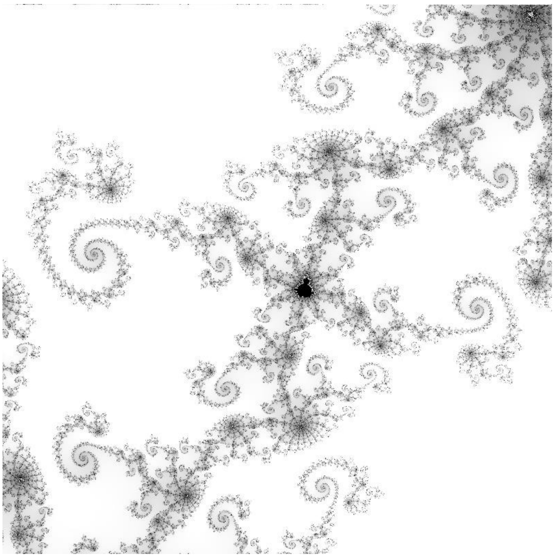

This is one of the numerous minibrots in the spiral arms with its dendrite running down the middle and two prominent spiral synapses on either side.

Looking on the other side of lobe 2 (i.e. between lobe 1:2:2 and lobe 1:2:3), the neurons frankly look a bit of a mess; but the tips of the dendrites are rather lovely. Here is the tip of dendrite number 8 (the largest one) on lobe 1:2:3›$2^6$ and a close up one of the minibrots which link the order-15 synapses:

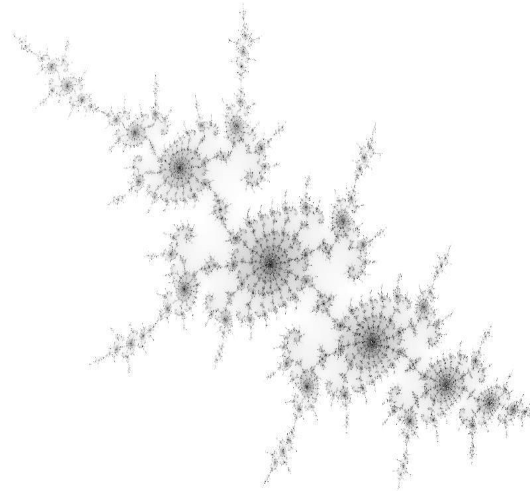

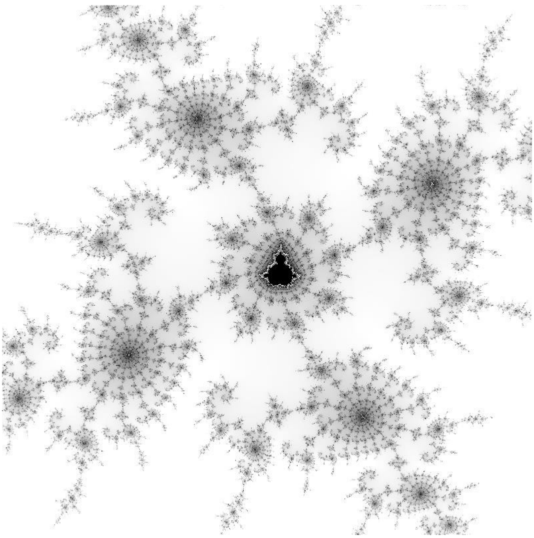

It is interesting to note that, while most dendrites curl round into infinite spirals, these are basically straight. Here is one of the dendrites on lobe 1:2:2:15 together with a close up of the minibrot just to the left of centre:

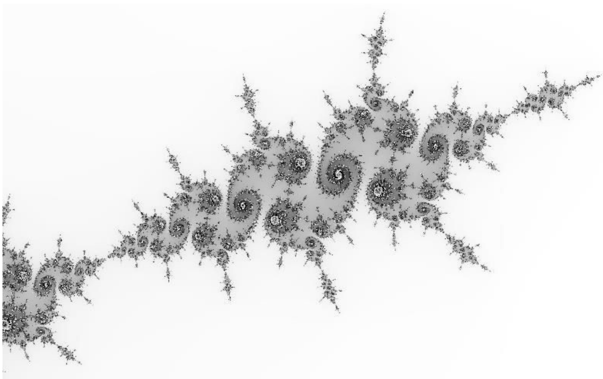

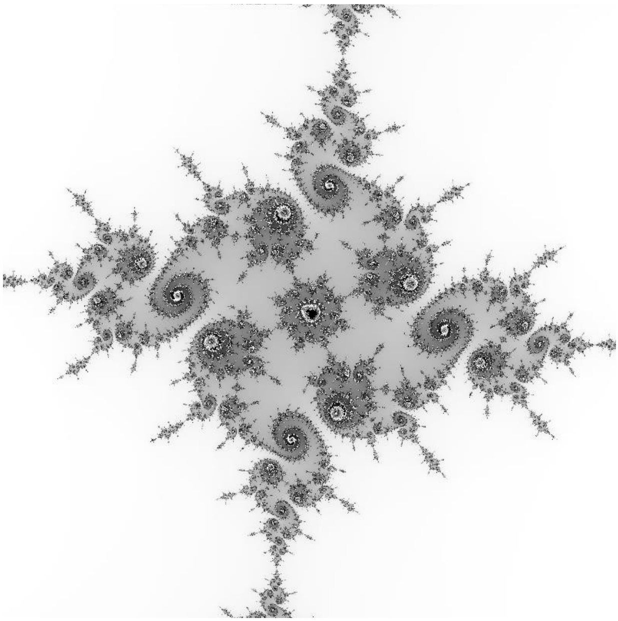

Other interesting dendrites can be found in places like 1:2:3:10 and many interesting patterns can be found within these structures which have been likened to sceptres.

# Neurons on the Other Major Lobes

We have seen that lobe 1:3 has order 3 but what does this mean in terms of its tertiary lobes? Here is an image of lobe 1:3:5:

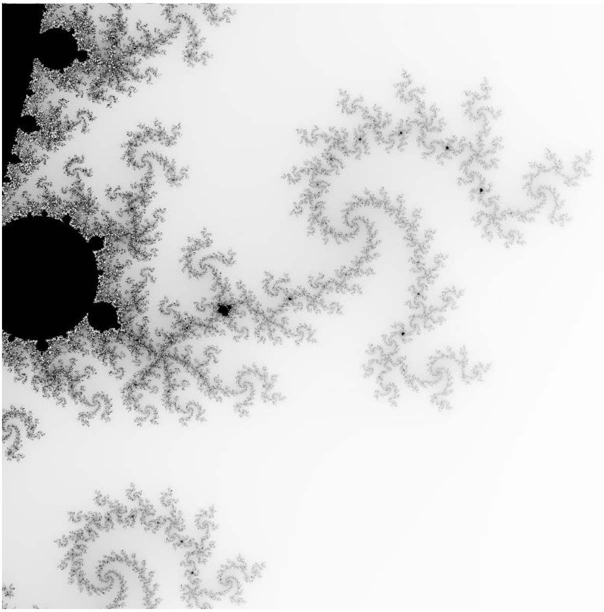

It is immediately obvious that we have synapses both of order 5 and order 3. The **principal synapse** is of order 5 and is attached to the lobe by a short straight axon but we can also see a beautiful triple spiral structure on the end of every dendrite which is, of course, a synapse of order 3. If you look at any of the dendrites in more detail you will find many more 'straight' order 5 synapses and spiral ones of order 3.

When we looked at the the secondary lobes of lobe 2 we saw double spirals. Here is the spiral synapse of lobe 1:3:10. It is still an order 3 synapse but it is given a lot more twists because of the high order of the lobe to which it is attached.

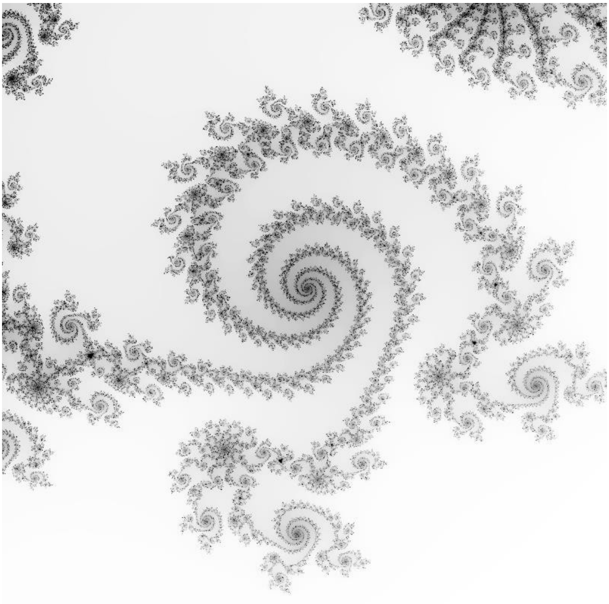

Now have a look at the main neuron of lobe 1:5:3:

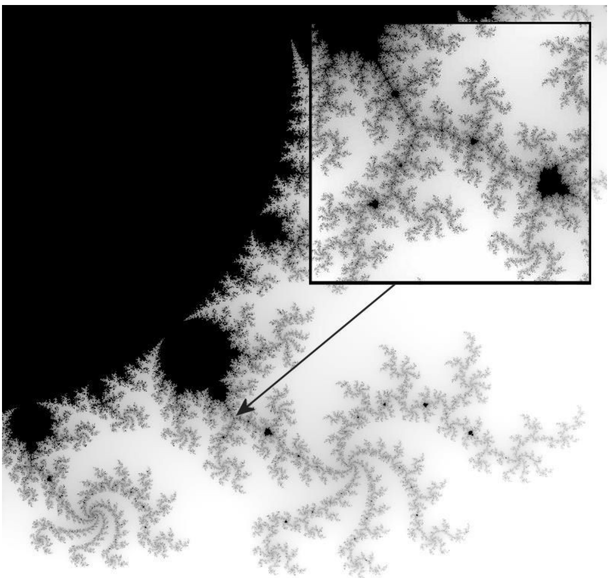

The order 5 spiral synapse is clearly visible but this is not the principal synapse which always lies between the lobe and the largest minibrot. It is pointed out in the illustration and has order 3. In fact you can find lots more order 3 synapses if you look for them.

We are now beginning to get a feel for the structure of a neuron. The synapses seem to mirror the periodicity of the lobes to which the neuron is attached.

But there is more.

Neurons do not actually sprout from a lobe such as 1:3:5 because there is always an infinite string of subsidiary lobes of order 2 between it and the place where the neuron actually starts. We could say that the principal neuron of lobe 3 actually starts from lobe 1:3:5:2:2:2:... We should therefore expect to find an infinite number of order 2 synapses as well as the synapses of order 3 and 5. And if the rule really holds strictly we should also find a synapse of order 1 somewhere. So where are they?

Well, the latter is easy to find. Every dendrite ends somewhere – and a terminus has only one branch so it is indeed a synapse of order 1. (You may recall that every synapse is a Misiurewicz point and these are the values of $c$ where the point $z$ wanders round for a while before hitting on a periodically self-mapping point.)

Now what about the order 2 synapses? You will recall that there is always an absolutely straight section between every lobe and the first major synapse. This is because there are an infinite number of order 2 lobes stacked one after another along the axon. In fact there are an infinite number of order 2 synapses between any pair of minibrots all along the dendrites as well.

Now here's a little test: using the principles which have been outlined so far, can you tell which lobes the following images are of?

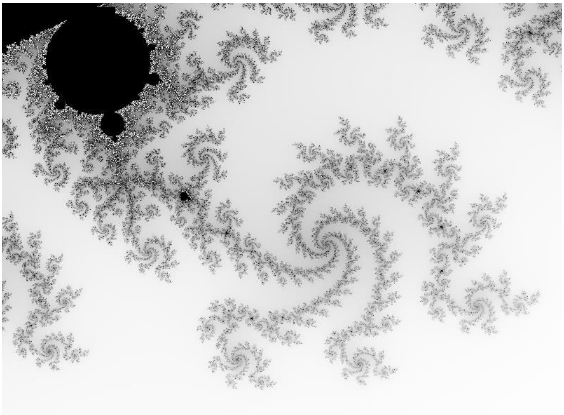

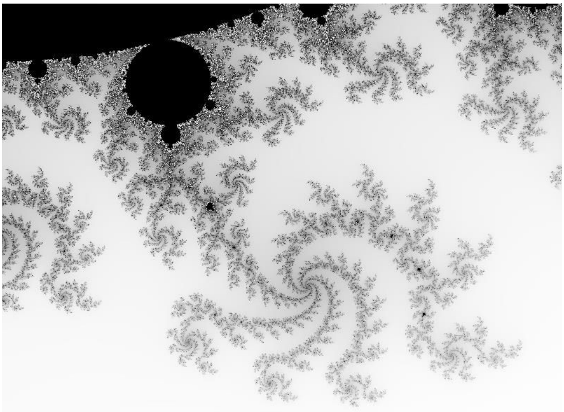

In each case the most obvious feature is spiral synapse but this is not really the most important feature. That is the first synapse at the end of the axon – the **principal synapse**. This is the order of the last lobe in the sequence (not counting all the order 2 lobes). The one at the top has order 5 and the one below order 4. Knowing the order of a lobe does not tell us uniquely which actual lobe it is because, as we have seen, different lobes can have the same order but I can tell you that these are both major lobes. (In fact we can infer this from the fact that in both cases the dendrites at the principal synapse form a single sequence descending in size.)[3]

3 The lobe labels are 1:4:5 (top) and 1:5:4 (bottom).

This is the principal neuron of lobe 1:3:2:4. Note how the synapses can be read backwards from the tip.

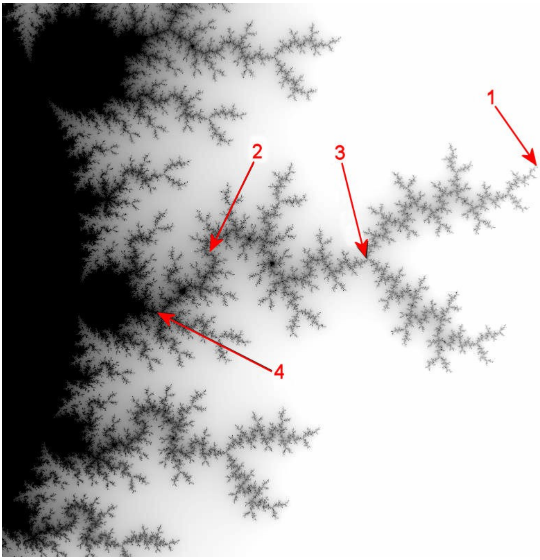


and this one is 1:2:3:6

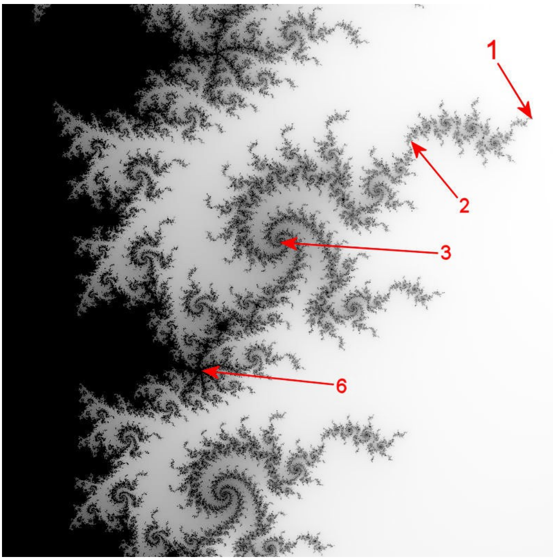

The following four neurons all have synapses of the same order nut some of them are from minor lobes. See if you can identify them. The answers are at the bottom.

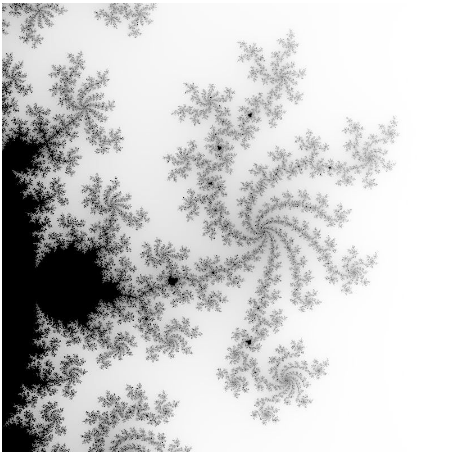
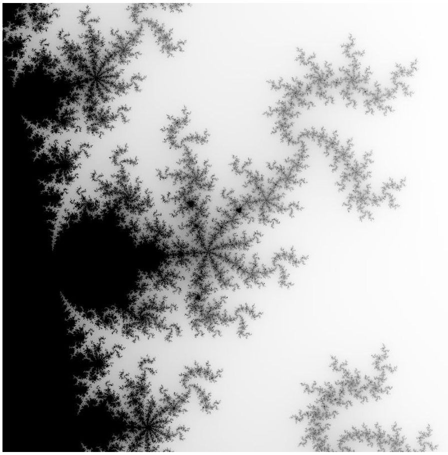
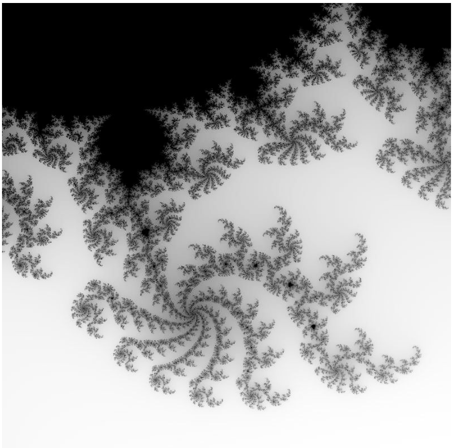
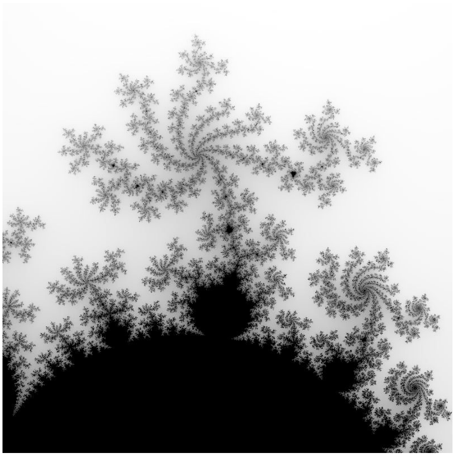

*Anticlockwise from the bottom left the answers are: 1:7:3, 1:3›2›2:3. 1:3:4›3 and 1:4›3:3*

# Axons and Dendrites

Every synapse has one axon and several dendrites. The axon is ultimately connected to the main lobe while the dendrites terminate in a synapse of order 1. You can tell the difference between an axon and a dendrite because the minibrots on the axon point *towards* the principal synapse while the minibrots on the dendrites point *away* from the synapse. You can see this clearly when the order is small but it is not always so easy to see the difference when the order of the lobe is large. Here is the principal synapse of lobe 13. Can you see which of the 13 branches is the axon?

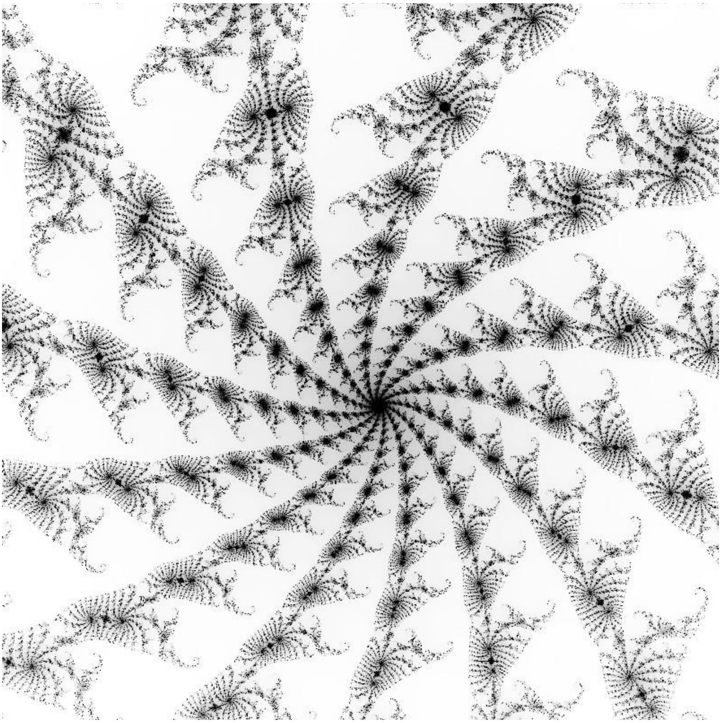

At first glance, all the branches look pretty much identical but a close up look at the two 'eyes' near the top of the image reveals subtle differences.

In the upper of the two images on this page, the axon flows straight down through the minibrot as indicated by the arrows but in the lower image, the 'axis' of the dendrite follows a zig-zag path up from the bottom and out at the top right hand corner.

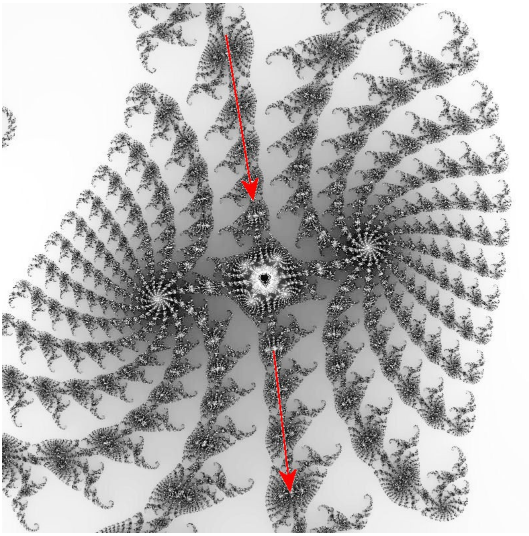

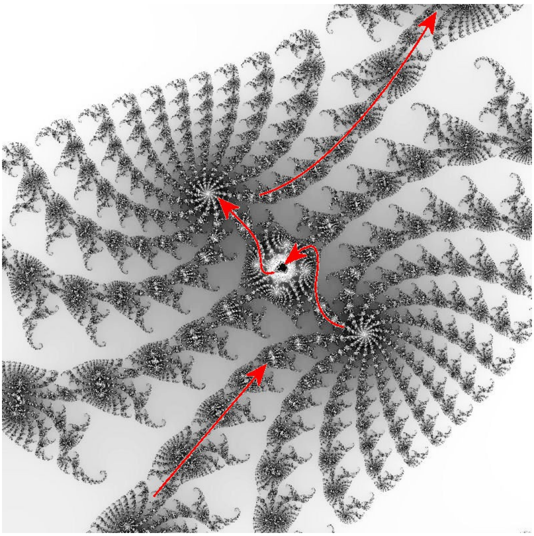

The upper image is therefore the axon and the lower one the dendrite.

# THE FRACTAL STRUCTURE OF A NEURON

So far we have been concentrating on the synapses but it is pretty obvious that between every pair of synapses there is at least one minibrot. So how many minibrots are there along a single neuron?

Lets clarify some of the ground rules which govern the structure of a neuron.

- Every neuron starts with a more or less straight section called the **principal axon** which leads to the **principal synapse** which has a total of $N$ branches (including the axon), $N$ being the **principal order** (or just 'order') of the lobe.
- The principal axon consists of an infinite series of minibrots and order 2 synapses.
- At the principal synapse, the principal axon branches into $N-1$ **dendrites** which have identical topology consisting of a sequence of synapses separated by minibrots.
- Between every pair of synapses on the same dendrite there are an infinite number of minibrots (and between every pair of minibrots there are an infinite number of synapses, many of which, being of order 2, are invisible).
- In the case of lobes with labels like $1{:}a{:}b{:}...z$ , the principal synapse has order $z$ and the subsequent synapses will appear in reverse order terminating with a synapse of order 1.

We have seen that every axon is, in fact, attached to an infinite series of order 2 lobes so that what we are calling lobe 3 (or more accurately 1:3) is in fact lobe 1:3:2:2:2:2... but let us ignore this for the moment and pretend that there was such a lobe as lobe 1:3. What would its neuron look like? My guess is that it would look something like this:

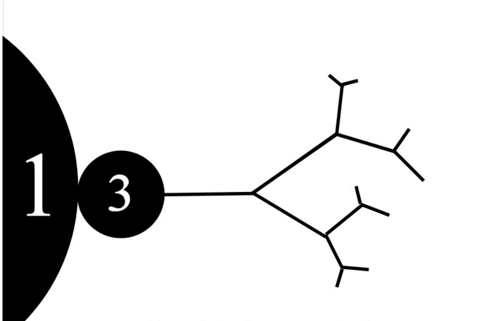


*Simplified neuron 1:3*

where all the dendrites branch into 3 an infinite number of times. Bearing in mind that the terminus is a synapse of order 1 we can list the series of synapses on the neuron as 3333...1 or more compactly {3}1 where the curly brackets indicate an infinite number of synapses.

What about neurons 1:3:5 and 1:5:3? My guesses would be:

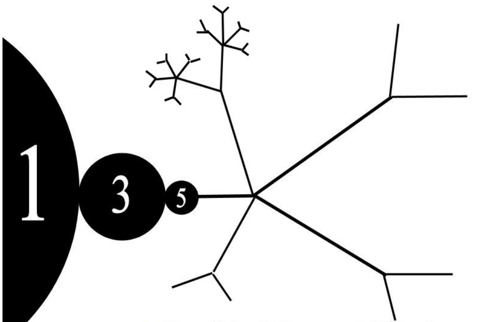


*Simplified Neuron 1:3:5*

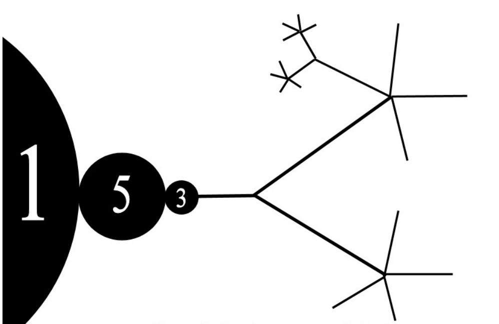


*Simplified neuron 1:5:3*

Note that between the principal synapse and the second synapse, there will be an infinite series of smaller synapses of the same order as the principal one (not shown). In other words the structure of the 1:3:5 neuron is:

((555...) 3) ((555...) 3) ((555...) 3) … 1

which we could write as:

$$\{\{5\}3\}1$$

The 1:5:3 neuron has the structure:

$$\{\{3\}5\}1$$

Since we have an infinite series of infinite series, loosely speaking the total number of synapses will be $\infty^2$ which equals $\infty$.

Now what happens when we add another lobe of order 2. Using the notation introduced above the 1:3:5:2 neuron will have the structure$\{\{\{2\}5\}3\}1$. Now since order 2 synapses are effectively invisible being just straight lines, the neuron will not look significantly different, but if we include the order 2 synapses, the total number of synapses will be $\infty^3$. If we add a second order 2 lobe (i.e. lobe 1:3:5:2:2) we have the fractal structure$\{\{\{\{2\}2\}5\}3\}1$ which contains $\infty^4$ synapses.

OK – so now lets go the whole hog. What structure does the neuron of lobe 1:3:5:2:2:2... have? I believe the answer will be:

$$\{\{\{...\{\{2\}2\}...\}5\}3\}1$$

which will contain $\infty^\infty$ synapses. Now $\infty^\infty$ is an uncountable number called $\aleph_1$. And if there are an uncountable number of synapses, there must be an uncountable number of minibrots too!

I have made two outrageous claims which seem to be contradictory. First I said that I believe that all synapses are Misiurewicz points. Now a Misiurewicz point is one which jumps around for *n* iterations and then enters a finite periodic cycle of period *p*. What this means is that $P_n$ must be equal to $P_{n+p}$ where $P_n$ is the $n^{\text{th}}$ iterate of the function $\boldsymbol{z'} = \boldsymbol{z}^2 + \boldsymbol{c}$ starting at $\boldsymbol{z} = 0$. To make this clearer, we shall list the first few iterations of the starting point $z = 0$:

$$
\begin{aligned}
P_0&:\ z_0 = 0\\
P_1&:\ z_1 = c\\
P_2&:\ z_2 = c^2 + c\\
P_3&:\ z_3 = (c^2 + c)^2 + c = c^4 + 2c^3 + c^2 + c\\
P_4&:\ z_4 = (c^4 + 2c^3 + c^2 + c)^2 + c
\end{aligned}
$$

and the Misiurewicz point $M_{2,1}$ (i.e. the value of c which produces a sequence with a pre-period of 2 and a period of 1) would be the root of the equation $P_3 = P_2$ which happens to be -2.

Now it is easy to see that if $c$ is a Misiurewicz point, it must be the root of a polynomial of finite order and hence must be a member of the countable algebraic numbers. This seems to be at odds with my claim that there are an uncountable number of order 2 synapses on any neuron. The solution to this puzzle is that my 'proof' that there are an uncountable number of synapses on the axon depends on there being an infinite number of order 2 lobes stacked on top of one another. This admits of the possibility that some of the synaptic points will be solutions to polynomials of *infinite* order. These points will not be algebraic and the pre-period will not be finite. Strictly speaking, these points are therefore not Misiurewicz points but I like to think of them as such – only with an infinite pre-period.

But here comes another difficulty. Minibrots, almost by definition have finite extent. They are in fact basins of attraction. Now you can easily divide a line into a infinite (i.e. countable) number of finite segments (take a line, remove the middle third excepting the ends, repeat …) but you cannot divide a line into a uncountable number of finite segments. So how can there be an uncountable number of minibrots?

The solution is the same. The extent of a minibrot will be a function of the depth of the lobe with which it is associated. If we allow an infinite number of order 2 lobes, the 'last' minibrot will have zero extent.

In the last analysis, synapses and minibrots become the same – truly chaotic points which never repeat.

# The Antenna

The primary neuron of lobe 2 – the one along the X axis, often called the 'antenna' – is, of course, special. It has order 2 so all the order 2 synapses lie along a straight line. In fact it is impossible to tell where the synapses are; all you can see is a chain of minibrots of varying sizes. Each of these minibrots is organised in the same way as the main one. The principal synapse is approximately half way between the main lobe and the largest minibrot on the axis and is shown with a red arrow in the image below.

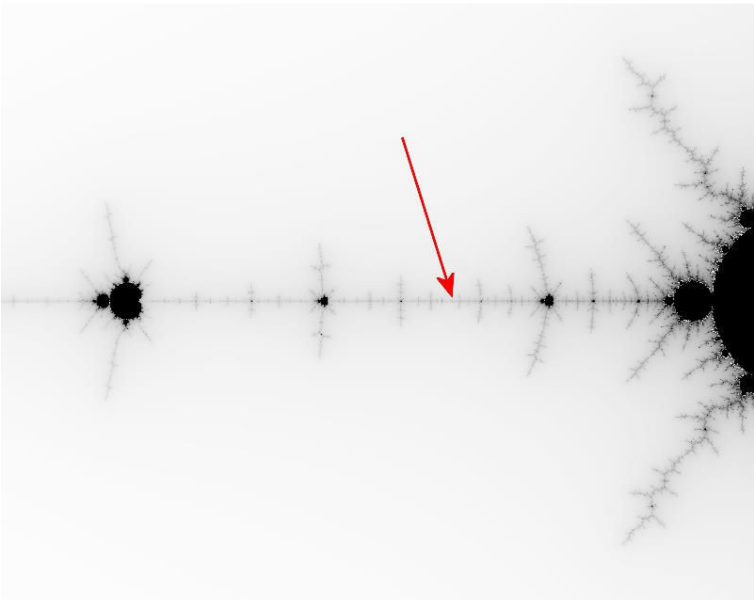

This point is in fact the Misiurewicz point $M_{3,1}$ i.e. it is one of the solutions to the equation $P_3 = P_4$. This equation reduces to $\boldsymbol{c}^3 + 2\boldsymbol{c}^2 + 2\boldsymbol{c} + 2 = 0$ whose only real solution is –1.543

In order to understand the structure of neurons better, it is useful to plot the behaviour of the point $z = 0$ for all values of $c$ from -2 to 0.25 along the main axon. This generates the following chaos map:

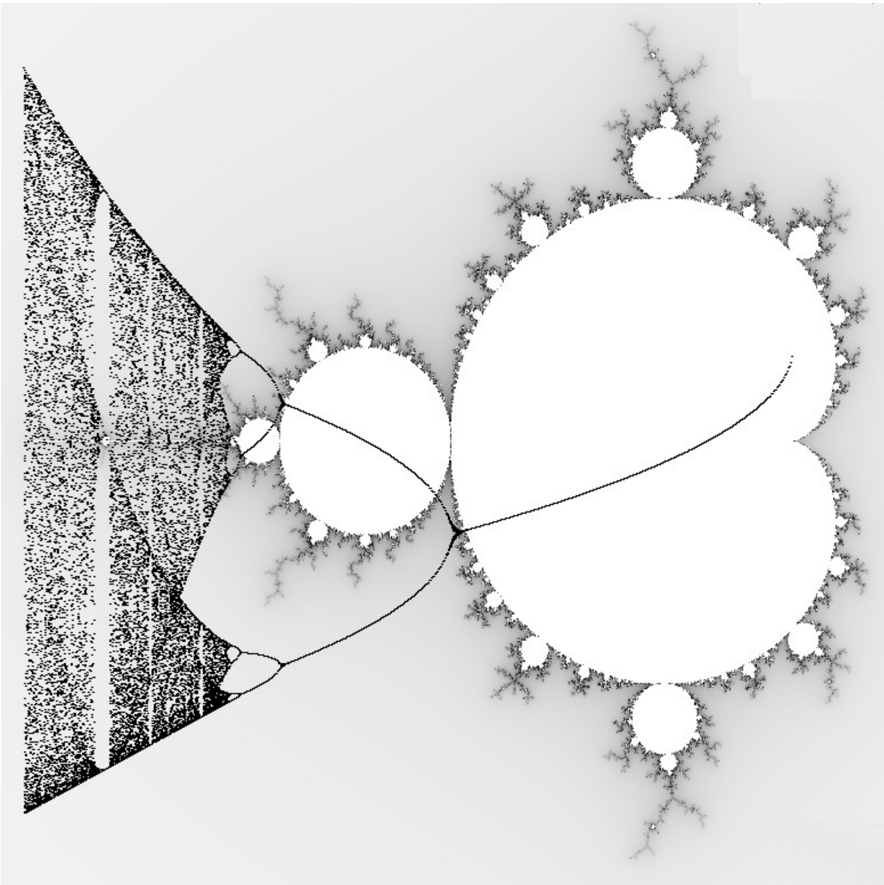

For anyone who has played with the equation $x' = \mathrm{A}x(1 - x)$ this will be very familiar. Reading it from right to left, we see that the point where lobe 1 meets lobe 2, the function bifurcates; and again and again through successive order 2 lobes. At a critical point (The Feigenbaum point $x$ = -1.4011551890… ) we reach the end of the series of order 2 lobes and the first order 2 synapse. Each subsequent island of stability in the chaos map indicates the position of a minibrot, the largest of which is clearly visible in the chaos map. (For a detailed explanation of why these basins of attraction exist and how to calculate their positions and periodicities, see the companion volume 'Chaos and the Logistic Equation'.)

Every value of $\boldsymbol{c}$ can be put into one of five categories:

**Intrinsically periodic** or **superstable** points**.** These are the values of $\boldsymbol{c}$ which cause $z$ to immediately enter a periodic orbit. These orbits therefore all contain the point $z = 0$.

There is only 1 point which maps onto itself and that is $\boldsymbol{c} = 0$.

There is only one point which immediately enters a cycle of period 2 which is -1. This is the centre of lobe 2. It is the other solution to the equation $P_2 = P_0 : \boldsymbol{c}^2 + \boldsymbol{c} = 0$ (apart from $\boldsymbol{c} = 0$ that is.)

There are three solutions to the equation $P_3 = P_\mathbf{0} : \boldsymbol{c}^4 + 2\boldsymbol{c}^3 + \boldsymbol{c}^2 + \boldsymbol{c} = 0$ namely $\boldsymbol{c} = –1.76$, $\boldsymbol{c} = (–0.12 + 0.75\mathbf{i})$ and $\boldsymbol{c} = (–0.12 – 0.75\mathrm{i})$. The first is in the centre of the largest minibrot on the axon and it is the most prominent 'island of stability' in the chaos map. The other two solutions lie in the centre of lobes 1:3 and 1:2›1. At once it becomes clear why these latter lobes have periodicity 3. It also means that the largest minibrot on the antenna has a basic periodicity of 3 too.

The next equation ($P_4 = P_\mathbf{0}$) is too long to write out and has 8 roots of which 2 are real, namely $\boldsymbol{c} = –1.32$ and $\boldsymbol{c} = –1.94$. The former is, of course, the centre of lobe 2:2 but surprisingly the latter is in one of the smaller minibrots near the end of the axon. Both of these have periodicity 4.

In general the equation $P_p = P_0$ will have $2^p$ roots, all of which will be at the centre of one of the lobes on one of the countless minibrots which litter the Mandelbrot map. Now, perhaps, you can see why there are so many of them!

These points are powerful attractors and all the points in the vicinity converge on them. (The reason for this can be found in the companion volume 'Chaos and the Logistic Equation') We might call these neighbouring points **asymptotically periodic**. They enter a cycle which never repeats but gets closer and closer to a periodic cycle.

Next on our list are the **Misiurewicz** points – the synapses. These are values of $\boldsymbol{c}$ where $\boldsymbol{z}$ jumps around for a while and then enter a periodic cycle. The point $M_{np}$ makes $n$ jumps before entering a cycle of period $p$.

As an example consider the case $c = -2$. This goes 0, –2, 2, 2, 2 … Its pre-period is 2 and its period is 1. It is therefore designated $M_{2,1}$. As we have seen, these points are unique solutions to a polynomial equation of finite order. Now in the companion volume 'Chaos and the Logistic Equation' I showed that Misiurewicz points can never lie inside a basin of attraction. What this means is that all Misiurewicz points are unstable. If $c$ is very close to the Misiurewicz point, $z$ will just miss the stable point and soon wander off to infinity. Since every Misiurewicz point is one of the roots of a finite equation, they are countable in number.

The fourth kind of point are the **Feigenbaum** points. These are the points which terminate the sequence of order 2 lobes which are to be found at the end of every minibrot. It is not possible to write down an equation for these points as they are limit values of an infinite series of equations. We can, however, say that, like the Misiurewicz points, they are countable in number because there are a countable number of minibrots.

This still leaves us with the possibility – indeed the certainty – that along any neuron there exist a huge number of values of $c$ which never cause $z$ to enter a periodic cycle but cause it to jump about chaotically for ever and ever.. I think we can reasonably call them **infinite Misiurewicz** points on the basis that they are really ordinary Misiurewicz points but with an infinite pre-period. And just as we associate the finite Misiurewicz points with a particular order 2 lobe on the axon, I think we can reasonably say that these infinite Misiurewicz points are synapses associated with the limiting order 2 lobe.

Alternatively you could argue that these points are simply the roots of the equation $P_p = P_0$ as $p \to \infty$. On this view they are super-stable but their period in infinite and their basin of attraction is infinitely small.

We know that these points must exist but I strongly suspect that it will prove impossible to calculate a single one of them!

# The Shape of the Lobes

We noted in the last chapter that the centres of all the lobes of periodicity $p$ are solutions of the equation $P_p = P_0 = 0$. The big question that now arises is this: what determines the size and shape of the region of stability round each intrinsically periodic point? Why does lobe 1 have the shape of a cardioid and why are all the smaller lobes circular?

For example, take the superstable centre of lobe 1:2. To calculate the superstable point we must solve the equation $P_2 = \boldsymbol{c}^2 + \boldsymbol{c} = 0$ and, as we have seen, the solutions are: $\boldsymbol{c} = 0$ and $\boldsymbol{c} = -1$. The first of these is, of course, the solution to the equation $P_1 = P_0$ (because any solution with period 1 also has period 2) but the second is the one we are looking for. With this value of $\boldsymbol{c}$ the point $\boldsymbol{z}$ cycles through the following points :

$$(0, 0) \rightarrow (-1, 0) \rightarrow (0, 0) \text{ etc.}$$

Now let us investigate how this cycle changes when we make small changes in the starting point (still using $\boldsymbol{c} = -1$).

$$f^2(\boldsymbol{z}, -1) = (\boldsymbol{z}^2 - 1)^2 - 1 = \boldsymbol{z}^4 - 2\boldsymbol{z}^2 \tag{1}$$

The graph of the real part of this function looks has a sort of double saddle-shape. (The colours indicate the magnitude of the imaginary component.)

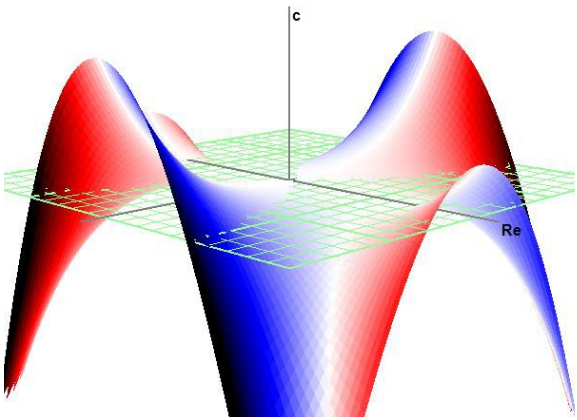

If we look at the vertical section along the real axis we see that it has two prominent penduline lobes at +1 and -1.

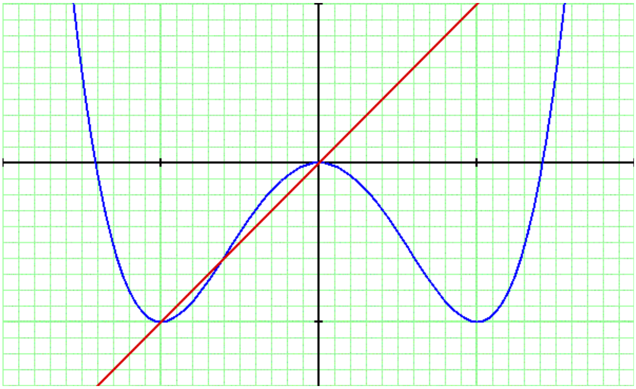

We shall be particularly interested in the gradient of this function and especially so at the points where $f^2(\boldsymbol{z}, \boldsymbol{c}) = \boldsymbol{z}$. – i.e. where the function is crossed by the 45° line. The gradient of the function is equal to $4\boldsymbol{z}^3 - 4\boldsymbol{z}$ or $4\boldsymbol{z}(\boldsymbol{z}^2 + \boldsymbol{c})$ and it is immediately clear it has zero gradient at both the origin and (when $\boldsymbol{c} = -1$) the point $\boldsymbol{z} = -1$. This is the reason why these two points are superstable.

Now in order to investigate how big lobe 2 is, we need to vary $\boldsymbol{c}$ in the region around (–1, 0). Suppose we move $\boldsymbol{c}$ to (–1.1, 0). This is the graph of the (real part of the) function

$$f^2(\boldsymbol{z}, -1.1) = (\boldsymbol{z}^2 - 1.1)^2 - 1.1 = \boldsymbol{z}^4 - 2.2\boldsymbol{z}^2 + 0.11 \qquad (2)$$

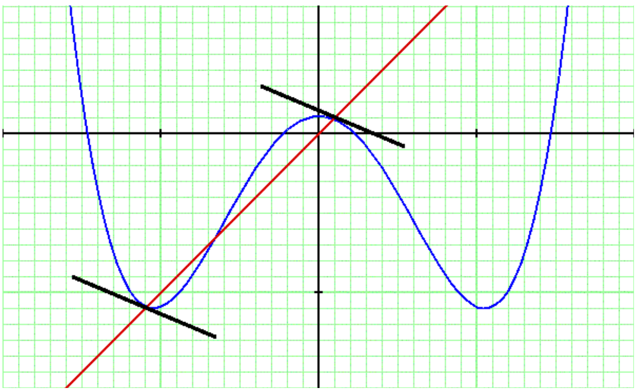

The penduline lobes drop a little further down and the central hump pokes above the axis. The solutions to the equation $f^2(\boldsymbol{z}, \boldsymbol{c}) = \boldsymbol{z}$ move slightly sideways and the gradients (shown in black) change. In my companion book 'Chaos and the Logistic Equation' I showed that, because of the iterative way in which the equations are constructed, these two gradients are always precisely equal. I also showed that, where solutions are stable, the absolute value of the gradient of the $f$ line must have a magnitude less than 1.

Now as $\boldsymbol{c}$ departs further and further from the superstable point, the $f$ curve becomes steeper and steeper and at a certain value, the gradients at the two solutions become equal to or greater than 1. At this point, the solutions become unstable and bifurcate. In the case of the $f^2$ Mandelbrot curve, this happens when $\boldsymbol{c} = -1.25$. Here is the graph of the function at this value of $\boldsymbol{c}$.:

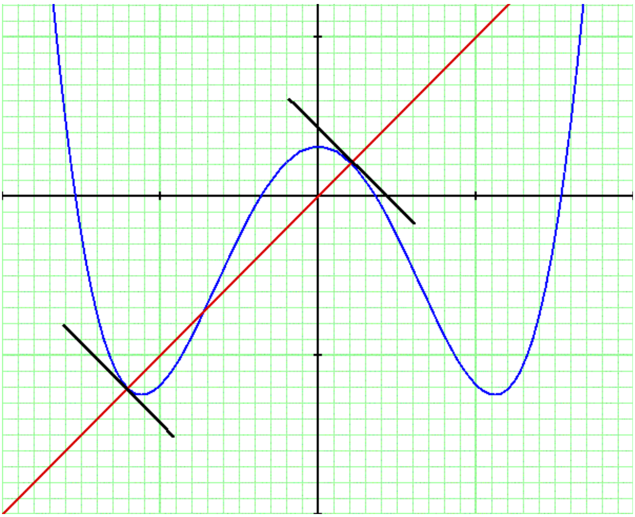

If instead of making $\boldsymbol{c}$, more negative we move it in the positive direction, a similar thing happens. The gradient becomes more and more positive until it reaches the point where the curve just touches the red line and the gradient = +1. In the case of the $f^2$ Mandelbrot curve, this happens when $\boldsymbol{c} = -0.75$. Here is the curve:

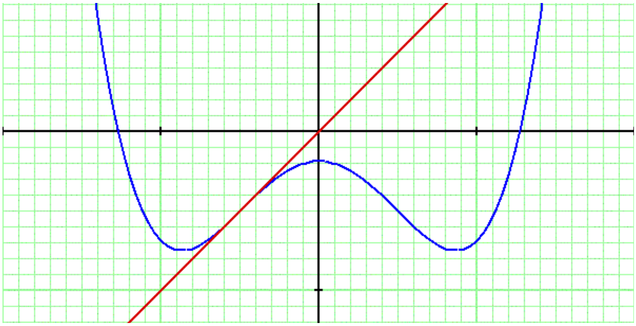

(In fact we observe that both the period 1 and period 2 solutions have merged which is why lobe 2 touches lobe 1 at the point (–0.75, 0))

To generalise what we have done, If we want to find the points where any lobe bifurcates – that is to say, if we want to find the boundary points of any lobe, we must find the locus of all the points $\boldsymbol{c}$ around the superstable point which are such that the modulus of the gradient of the function $f^{\mathrm{n}}(z, \boldsymbol{c})$ with respect to $z$ must be equal to 1 at the solutions of the equation $f^{\mathrm{n}}(z, \boldsymbol{c}) = z$. To be precise, in the case of lobe 2 we must find the locus of all the values of $\boldsymbol{c}$ which satisfy these two conditions:

$$z^4 + 2\boldsymbol{c}z^2 + \boldsymbol{c}^2 + \boldsymbol{c} = z \tag{3}$$

and

$$|4z^3 + 4\boldsymbol{c}z| = 1 \tag{4}$$

This is not an easy task to carry out analytically and it seems to me to be little short of miraculous that, in spite of the complex way in which it is generated, lobe 1:2 is, in fact, perfectly circular[4]. Equally miraculous is that, notwithstanding appearances, numerical investigations have shown that *none* of the other lobes are exactly circular!

Nevertheless, we can now appreciate why all the lobes (except lobe 1) are at least approximately circular. Here is a 3D picture of the modulus of the gradient of the $f^2$ function – i.e. $|4z^3 + 4\boldsymbol{c}z|$ when c = –1:

4 I have, however, managed to devise a fairly simple proof which is given in the appendix.

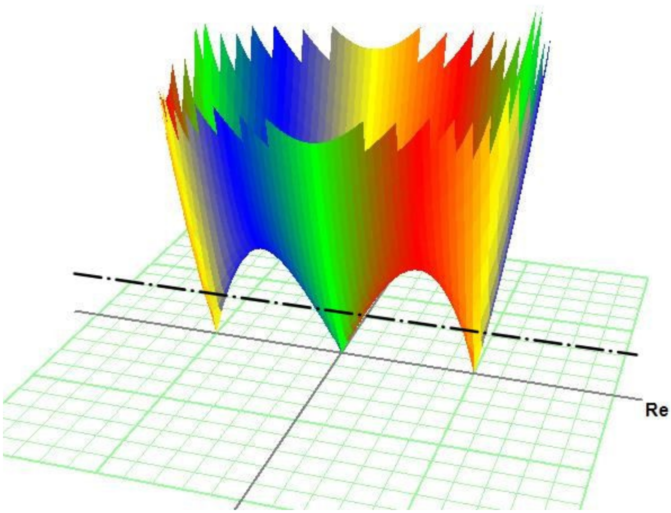


Notice how sharp the cones are which dip down to zero at $\boldsymbol{c} = -1$, 0 and +1. The height of the 'saddle' between the points is 2.25 so if we were to cut the graph off at a height of 1 (indicated roughly by the height of the dotted line above the real axis) the locus of points at this height around $\boldsymbol{c} = -1$ would be approximately circular.

Asimilar graph showing the gradient of the $f^3$ function around its superstable point would show and even more sharply pointed cones. I strongly suspect that of all the lobes with periodicity greater than 2, lobe 1:3 will be the most distorted and even that has an eccentricity of less than 1%.

But what about lobe 1? Surely we must be able to do the calculations for the simplest lobe!

Here $f^1(\boldsymbol{z}, \boldsymbol{c}) = \boldsymbol{z}^2 + \boldsymbol{c}$ and the self-mapping points are the solutions to the equation $\boldsymbol{z}^2 + \boldsymbol{c} = \boldsymbol{z}$ or $\boldsymbol{c} = \boldsymbol{z} - \boldsymbol{z}^2$.

Now the gradient of the function $f^1(\boldsymbol{z}, \boldsymbol{c}) = \boldsymbol{z}^2 + \boldsymbol{c}$ is just $2\boldsymbol{z}$ and we know that at the boundary of the lobe, the modulus of the gradient must be equal to 1. The fact that the gradient does not depend on $\boldsymbol{c}$ is crucial because it allows us to say that, whatever the value of $\boldsymbol{c}$, then modulus of $z$ must be 0.5.

We now have to find the locus of all the points $c$ such that $c = z - z^2$ where $|z| = 0.5$. Substituting $z = x + iy$ and trying to solve for $c$ leads to all sorts of difficulties (I have tried it!) but a geometrical argument reveals the answer easily.

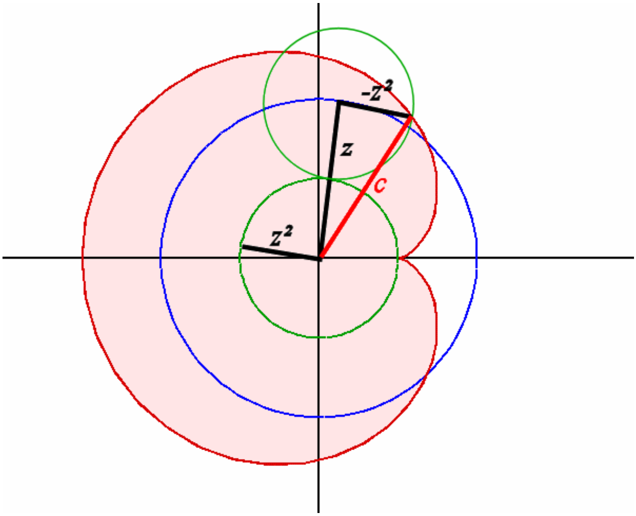


The blue circle has radius 0.5. $z$ must lie on this circle.

Squaring a complex number involves squaring the modulus and doubling the argument. $z^2$ therefore has a radius of 0.25 and rotates round the green circle at twice the rate of $z$.

$c$ is equal to the sum of $z$ and $-z^2$. Imagine the outer green circle rolling around the inner one carrying $z$ and $-z^2$ with it as if they were two linked rods. The free end will describe a cardioid and it is easy to see that the cusp of the cardioid will be at the point (0.25, 0) and the opposite pole at (–0.75, 0) which is exactly describes the shape and position of lobe 1!

# The Position and Size of the Lobes

We now know, at least in principle, why the lobes have the shape that they do. But what determines their size and position and why are the lobes so closely associated with the Stern-Brocot sequence?

If you look closely at the tertiary lobes on lobe 2 (illustrated below) you will see that the position of the lobe around the perimeter corresponds exactly with the rotation number of the lobe – that is to say, lobe 1:2:3 is 1 third of the way round the perimeter, lobe 1:2:3›2 is 2 fifths of the way round etc.[5]

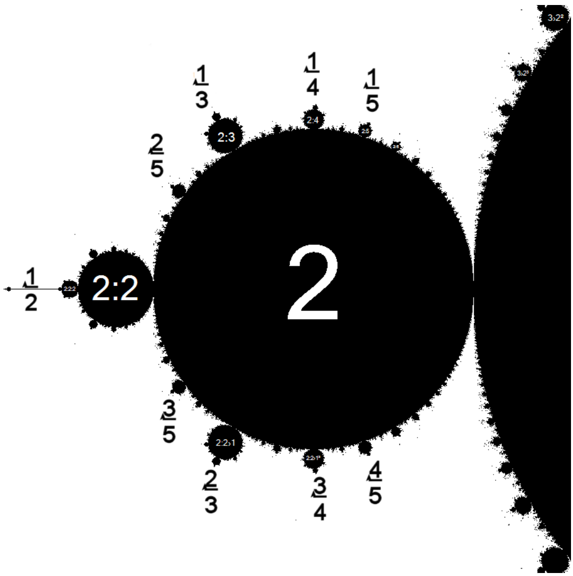


In addition, it appears that the diameter of the lobes is closely related to the periodicity – the smaller the periodicity, the larger the diameter and all lobes with the same periodicity have approximately the same diameter.

The fundamental question is – why? What is the deep reason for all this order and regularity?

5 I have not seen a proof of this statement but neither have I seen any evidence to contradict it.

The frank answer is – I do not know. Of course, there is a reason and, no doubt, any competent mathematician could prove that, for example, the lobe 1/3rd of the way round the lobe will have a periodicity of 3 – but this is not quite the same as explaining *why* this simple relation holds.

On page 51 I explained the reason why bifurcation occurs between lobe 1 and lobe 2 when the gradient of the $f$ function reaches $\pm 1$. There we restricted ourselves to the real axis. If, however, we consider the behaviour of $\boldsymbol{z}$ in the complex plane, we can find that 'bifircation' is not the only option. In fact, every point on the perimeter of any lobe is a 'fircation' point. You can see how this happens by tracing 100 or so iteration of $\boldsymbol{z}$ when $\boldsymbol{c}$ is very close to the edge of a lobe and near to one of the major lobes.

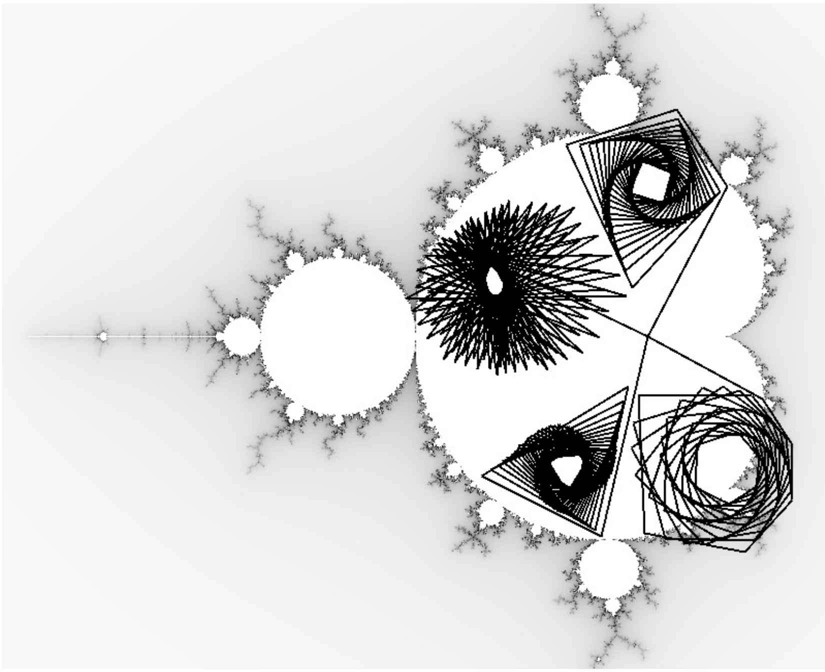

In each case you can see how $\boldsymbol{z}$ nearly becomes periodic but not quite. When $\boldsymbol{c}$ is close to lobe 4, $\boldsymbol{z}$ tries to home in on a square but each square is slightly skewed. It is clear that, in the limit, $\boldsymbol{z}$ repeatedly takes a step, then turns through an angle (basically equal to the rotation number). If the angle is an exact fraction (*q*/*p*) of a circle, the result will be a polygon with *p* sides and step size *q*! This explains why the periodicity and step size is closely related to the position of the lobe and, because of the special properties of the Stern-Brocot tree, why the lobes can be organised into sequences whose rotation numbers obey Farey addition.

The only thing that remains to be explained is why the diameter of the lobes is inversely related to the periodicity.

The illustration below maps out the solutions to the equations $P_2 = 0$, $P_3 = 0$, $P_4 = 0$ and some of the solutions to $P_5 = 0$::

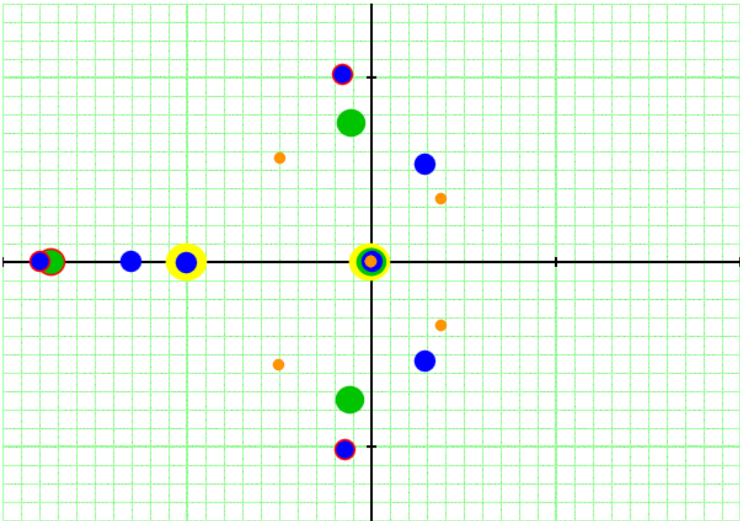

There are 2 solutions of periodicity 2 (shown in yellow), 4 of periodicity 3 (green), 7 of periodicity 4 (blue) and 15 of periodicity 5 (orange) of which 5 are shown. (The solutions with a red ring are not lobe centres, they are the centres of minibrots. Obviously the larger the periodicity, the greater the number of solutions and hence the greater the number of lobes with that periodicity. Of necessity, then, the larger the periodicity the smaller the lobes will have to be.

A remarkable example of this is illustrated by the construction known as Ford circles.

Draw a base line and on it draw two touching unit circles.

Now draw a the largest circle possible in the interstice. You will find that it has a diameter ¼ of the diameter of the unit circles and, of course, it is positioned half way between them.

Now draw two more circles in the interstices; amazingly these circles have a diameter of 1/9 and are positioned at 1/3 and 2/3.

When you come to do the next circles, you will find that the four interstices are not all the same size and there are only 2 places where you can fir a circle of diameter 1/16. But, remarkably, and in spite of the fact that the interstices are not all the same shape, there are *four* places where you can fit a circle of diameter 1/25. In fact, every time you come to a prime number $p$ the circles seem to know that all the interstices must be exactly the right shape to admit the addition of $p - 1$ circles of diameter $1/p^2$. [6]

Carry on ad infinitum. This is the astonishing result:

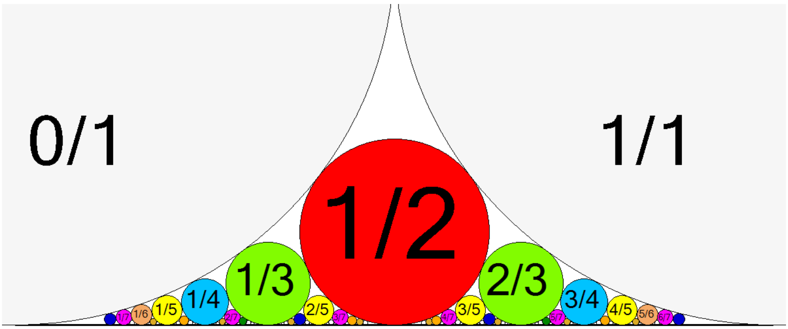


Look closely at the fractions which I have written in the circles. Each fraction is the Farey sum of the two circles above it and the whole nest of circles forms a Stern-Brocot tree!

Remove the two unit circles and bend the whole thing round a circle and this is what you get:

6 Many proofs of this theorem can be found on the internet.

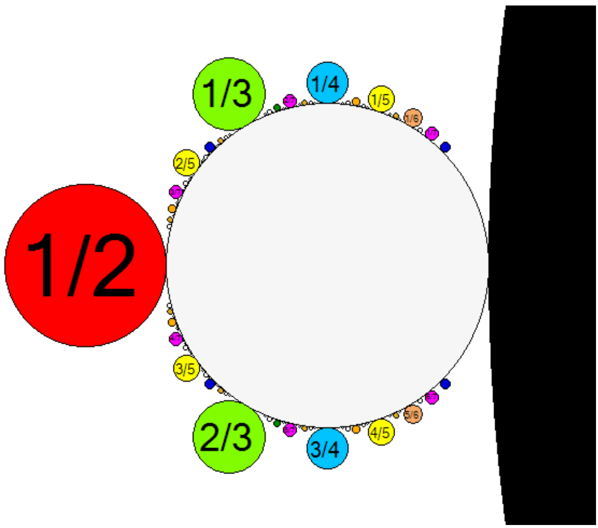


a Mandelbrot lobe complete with all its tertiary lobes!

One last question remains. Why are the angles of the secondary lobes on lobe 1 not equal to the rotation number? Lobe 3 is not 1 third of the way round – it is at the top, more like one quarter of the way round.

Well, as we have seem, lobe 1 is a bit of an exception. It is a cardioid, not a circle. When we bend the Ford circles round a cardioid, we have to distort the distances to accommodate. Look at the diagram opposite. The rotation number of the lobe determines the angle of $\boldsymbol{z}$ not $\boldsymbol{c}$.[7] As $\boldsymbol{z}$ rotates round from 0 to ½, $\boldsymbol{c}$ lags behind. When the rotation number is 1/3, $\boldsymbol{z}$ is at 120° to the axis but the angle which $\boldsymbol{c}$ makes is only a little over 90° putting lobe 3 nearly (but not quite) at the top.

---

7 Again, I have not seen a proof of this statement but I have no reason to suppose that it is not true. On the other hand, nothing should be taken as true in mathematics until it is formally proved so it should be born in mind that all my statements regarding the positions of the lobes around the perimeter of lobes 1 and 2 may not be absolutely accurate.

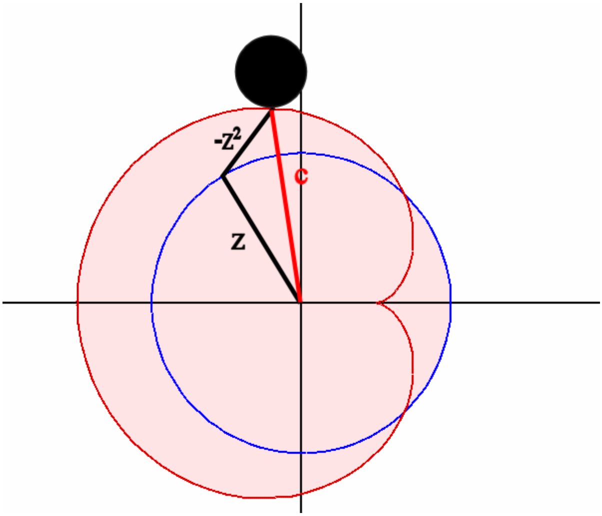


So every point on the boundary of a lobe which is associated with a rational fraction of the whole is a 'fircation' point and has a lobe attached whose periodicity is equal to the denominator of the fraction.

All well and good. But what about the irrational points? Or even the transcendental points? What happens there?

The simple answer is that when ***c*** equals one of these points, ***z*** wanders about chaotically. I shall call these points **aperiodic fircation** points. Obviously there are an uncountable number of them around the perimeter of every lobe.

On page 47 I discussed the possibility that on every neuron there were **infinite Misiurewicz** points which were similarly aperiodic but which were impossible to calculate. **Aperiodic fircation** points are similar but unlike the Misiurwicz points, we can easily name some of them. Those on the cardioid are of the form $\cos(t) - \cos(2t) + \mathbf{i}\,(\sin(t) - \sin(2t)$ where $t$ is any irrational number. Likewise all the points on the perimeter of lobe 2 have the form $1 - 0.25\cos(t) + \mathbf{i}\sin(t)$ and all those points where $t$ is irrational will be aperiodic[8].

8 But see the footnote on the previous page.

# Minibrots

It was noted in the previous chapter that there are potentially $2^{p-1}$ solutions to the equation $P_p = 0$ (although many of them will be duplicated or degenerate). All of these solutions will be centres of attraction but there are far more of them than there are lobes on the main brot. The other solutions lie at the centres of the numerous minibrots which litter the map.

The largest minibrot on the antenna has period 3. There are three minibrots of period 4: one on the antenna and two others in the neurons attached to lobes 1:3 and 1:2›1.

I believe that there are 15 different solutions of periodicity 5 of which 5 are accounted for on the main brot including the one at the origin (See the diagram on page 56). That leaves 10 minibrots. Two of then are on the antenna at $\boldsymbol{c}$ = – 1.626 and – 1.861. I challenge you to find the other 8!

Naturally, the periodicity of any lobe on a minibrot will be the product of its basic periodicity and the periodicity of the minibrot. This means that no lobe on a minibrot can have a periodicity $p$ which is a prime number. Alternatively we can say that the only basins of attraction with prime periodicities are either secondary lobes on the main cardioid or the main cardioid of a minibrot.

Now what can we say about the periodicity of all the other minibrots on a neuron? First we shall consider the *principal dendritic minibrot* i.e. the largest minibrot beyond beyond the principal synapse on each of the dendrites and also some of the other dendritic minibrots which lie between the principal one and the synapse.

Here is a short list with their periodicities (in clockwise order):

| Lobe | Periodicity of main lobe | Periodicity of the principal dendritic minibrot on each dendrite | Periodicity of the largest dendritic minibrot between the principal dendritic minibrot and the synapse |
|---|---|---|---|
| 1:2 | 2 | 3 | 5 |
| 1:3 | 3 | 4, 5 | 7, 8 |
| 1:4 | 4 | 5, 6, 7 | 9, 10, 11 |
| 1:5 | 5 | 6, 7, 8, 9 | 11, 12, 13, 14 |
| 1:3›2 | 5 | 8, 6, 9, 7 | 13, 11, 14, 12 |

The pattern is pretty clear. The periodicities of the largest minibrot simply carry on from the periodicity of the main lobe. Not surprisingly, the order of the minibrots in the last case steps round 2 at a time.

The last column lists the periodicities of the next largest minibrot as you step back along the dendrite towards the principal synapse. Its periodicity increases by the order of the synapse. Indeed, between the period 5 minibrot on the neuron of lobe 4 and its principal synapse you will find minibrots of periodicity 9, 13, 17 etc.

Now turn your attention to the minibrots beyond the principal dendritic minibrot and you will discover another remarkable sequence. The periodicities increase by 1 each time. Why should this be? I believe it is because we are heading towards a synapse of order 1 – the terminal synapse. Indeed, I strongly suspect that, starting at any dendritic minibrot of periodicity $p$ in one of the neurons of the secondary lobes and moving towards a synapse of order $o$ you will find a sequence of minibrots of periodicity $p + no$.

Surprisingly, this rule does not hold for dendritic minibrots on tertiary lobes. For example, consider the neuron attached to lobe 1:2:3.

It should be no surprise to find that the periodicity of its two principal dendritic minibrots are 8 and 10 and that the periodicities of the sequence of minibrots between them and the principal synapse increases by 6 each step.

Now if the rule held for minibrots beyond the principal dendritic minibrot, we would expect to find minibrots of periodicity 10, 12, 14 etc. but, unexpectedly, one of these minibrots has periodicity 5. (If you were having trouble finding the last pair of minibrots with periodicity 5 then here they are!)

What do these minibrots look like and how can we determine their periodicities? Below are images of the principal dendritic minibrots in the neurons of lobes 1:3 and 1:5. The most obvious thing about them is that, in addition to the synapses characteristic of the lobe to which the neuron is attached, every neuron has multiple synapses whose order is equal to the order of the neuron on which the minibrot is situated.

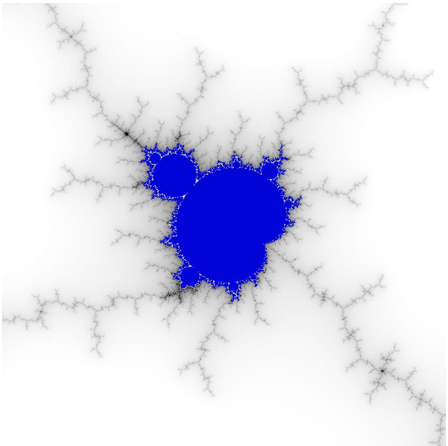

*1:3*

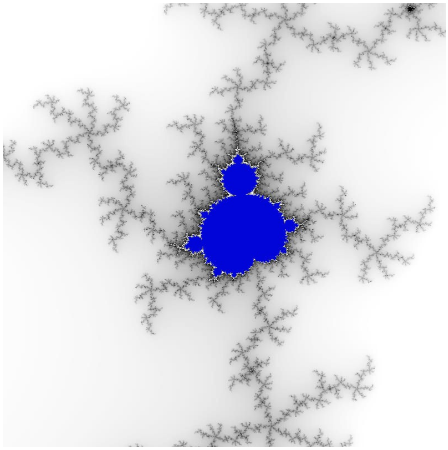

*1:5*

Unfortunately, there does not seem to be a way of determining the periodicity of the minibrot just by examining its neurons. The only way to do it is to trace the orbit of $z$ for a while and then cont the number of iterations before it returns close to a previous value. (This is the algorithm used by the author's program 'Mandelbrot Explorer'.)

This just leaves the minibrots on the axon. The largest is called the *principal axonal minibrot*.

Here are images of the principal axonal minibrots on the axons of lobes 1:2, 1:3, 1:5 and 1:3›2:

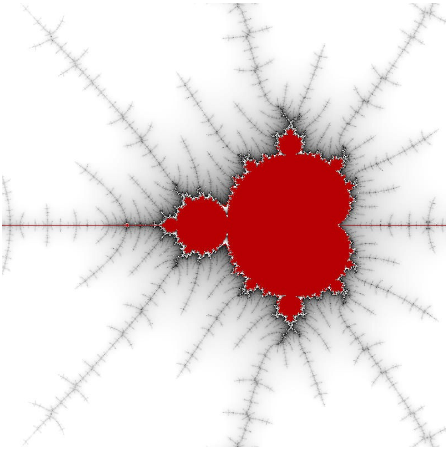

*1:2*

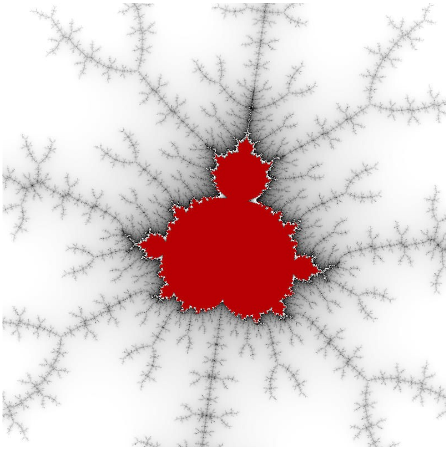

*1:3*

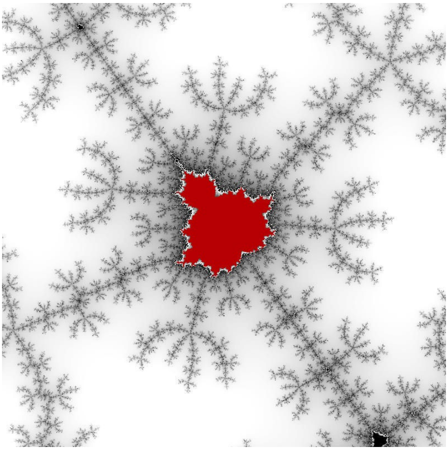

*1:3›2*

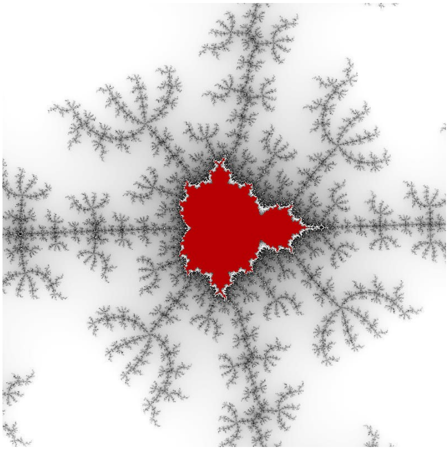

*1:5*

You will immediately notice that every neuron has a prominent synapse equal to the order of the main lobe to which the minibrot is linked. As we have seen this tells us the order of this lobe but it does not tell us the periodicity of the minibrot.

It turns out that the periodicity of these four minibrots is 6, 9, 15 and 15 respectively. Further investigation confirms that the periodicity of the principal axonal minibrot is always *3 times the order of the lobe* to which it is attached.

(Note that the principal axonal minibrot on the antenna is *not* the largest minibrot. The largest minibrot on the axon is the principal *dendritic* minibrot and has periodicity 3 (= 2 + 1); the former has periodicity 6 (= 2 × 3)) .

As you would expect, as we move towards the principal synapse, you will find a sequence of minibrots whose periodicities increase in steps equal to the order of the synapse.

Moving in the other direction (i.e. towards the lobe itself) the periodicities are always multiples of the periodicity of the lobe but I cannot discern an obvious pattern in their order.

In summary, here is a diagrammatic list of the periodicities of some of the larger minibrots in the neuron attached to lobe 1:3.

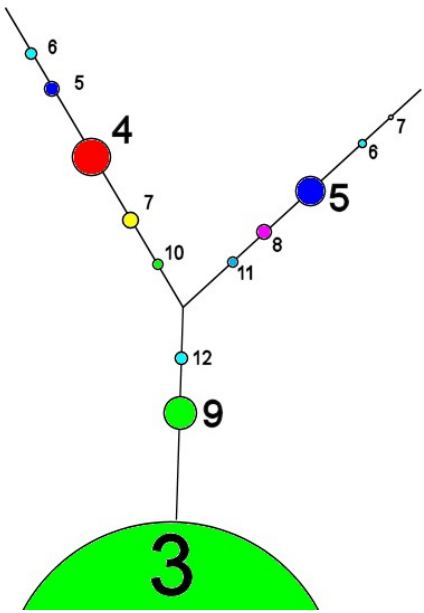

# Principal Synapses

That's enough theory: let's have some fun. But where shall we start looking for some interesting features? Let's start with the principal synapses. Here are close views of the principal synapse of lobes 6 and $3›2^9$ which we might call the 'snowflake synapse' and the 'peacock's tail synapse'.

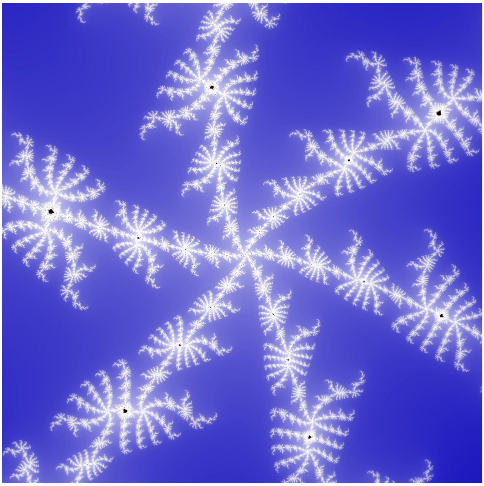

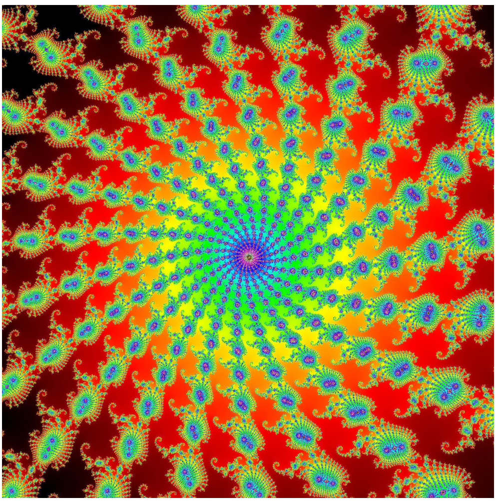

All synapses are self-similar – that is to say, they look exactly the same however much you magnify them – because they are Misiurewicz points. You will, however, notice that on the arms of the synapse, there appear to be synapses with twice as many arms. These are not real synapses though and if you magnify them sufficiently you will find a minibrot at the bottom with two genuine synapses, one on each side. A simple order 3 example is shown on page 25. Here is a more complex example. Can you identify which lobe it is on?[9]

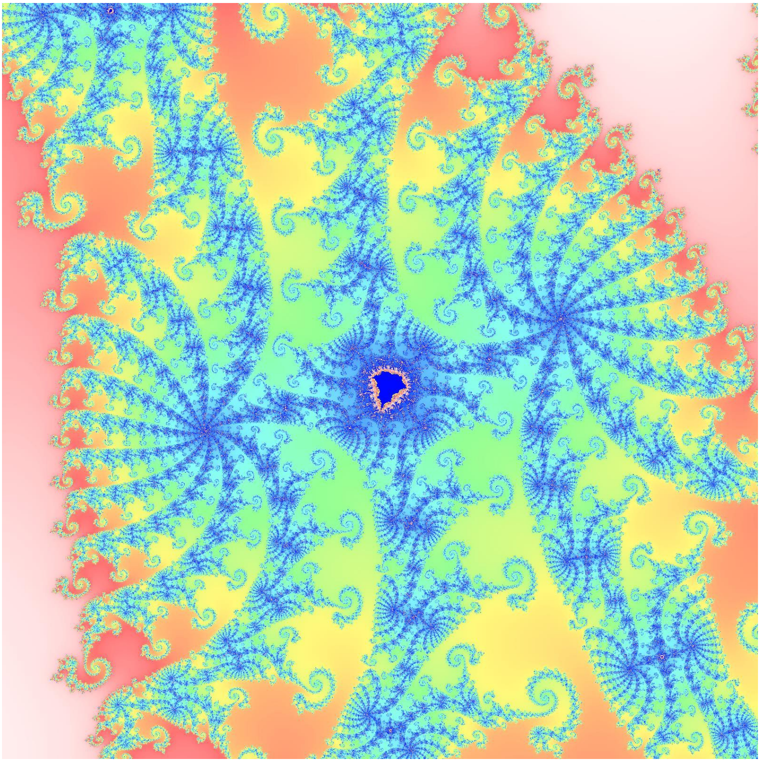

The main neuron always passes through the minibrot along its axis and there are 2 synapses on each side. If the order of the lobe is $N$, each of these synapses will contribute $N - 1$ dendrites which, together with the ingoing and outgoing axon make a total of $2(N - 1) + 2 = 2N$ arms. We can now see why, from a distance, this can look like a synapse with

9 Lobe 2:11

twice as many arms.

The details of these 'eyes' are infinitely varied and the source of many published images. Here is a close up of one of the 'eyes' of the principal synapse of lobe 3›$2^9$.

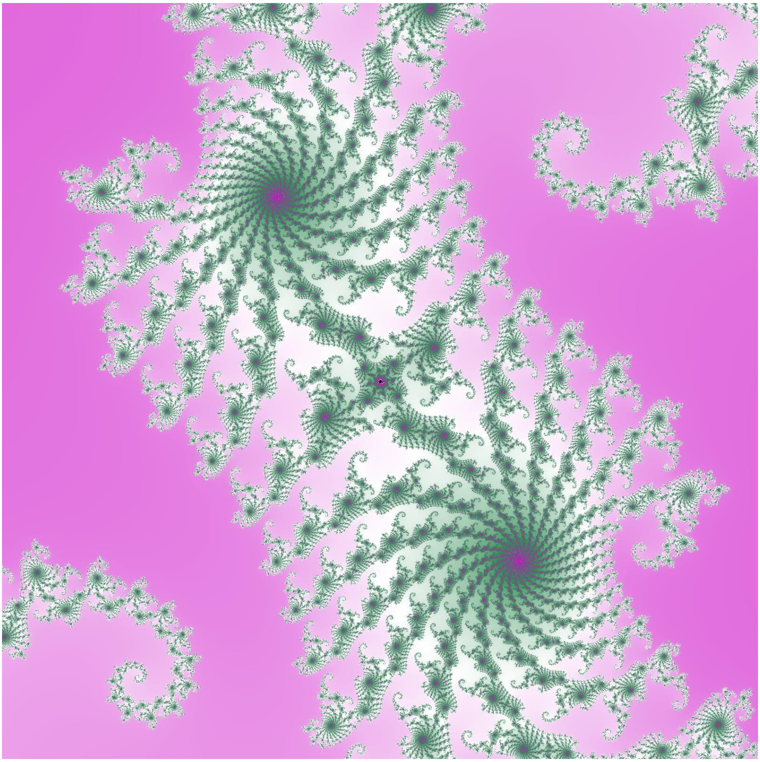

You can tell that the eye is on the main axon because the main axis passes straight through the minibrot. You can tell that the lobe is a primary lobe because the dendrites terminate in single spirals. You can tell that the order of the lobe is 21 because of the number of branches on each of the side synapses and you can tell that it is not lobe 21 because the dendrites form two groups, but whether it is on lobe 3›$2^9$ or lobe 5›$4^4$ which also has order 21 is more difficult to determine.

# Spiral Synapses

To find a spiral synapses with, say, 6 arms we need to look at the tertiary lobes of lobe number 6. This one is on lobe 1:6:5. The first image is a smooth image while the second is exactly the same location but alternate levels are picked out in green. The smooth areas between the spiral arms in the first image are revealed as huge roots which spiral down into infinite depths.

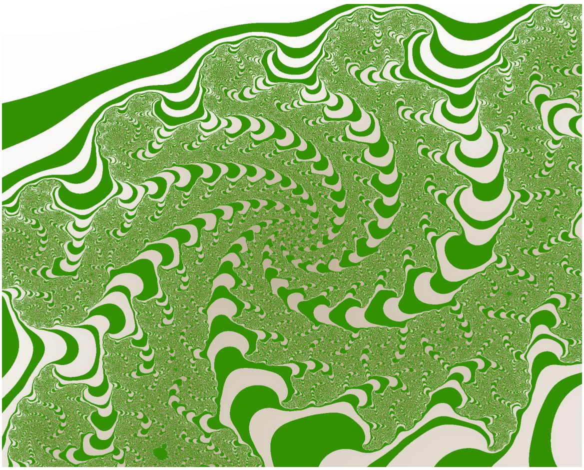

Compare that with the spiral synapse on lobe 1:2:6:5 which has order 2 spiral synapses tacked onto every branch.

As with principal synapses, it is the 'eyes' of the synapse which hold the greatest interest. Typically, each eye is bounded on each side by a spiral synapse of its own. The eyes increase in complexity as you go further down the main spiral.

Here is the spiral synapse of lobe 1:3:5 and a close up of one of its eyes.

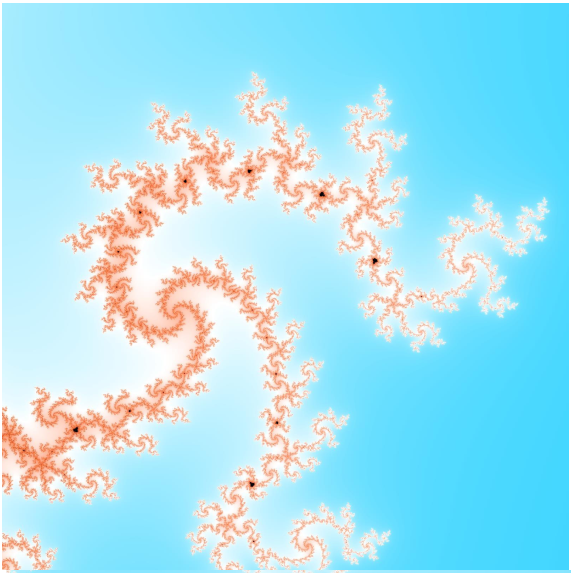

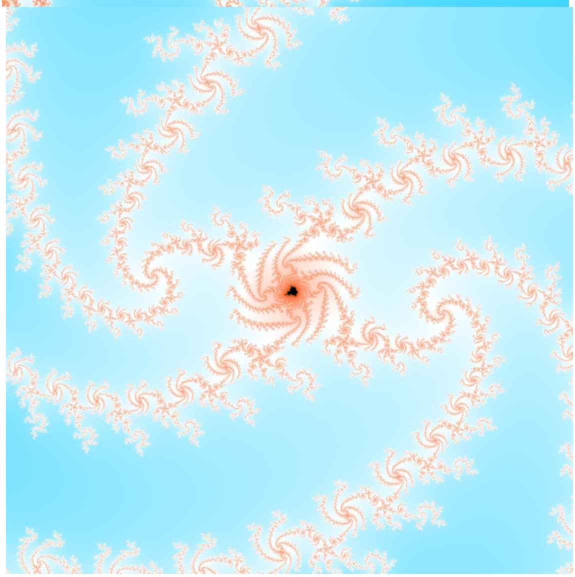

This is a pair of images from lobe 1:2:11. The 'eyes' here have lovely side spirals of order 2:

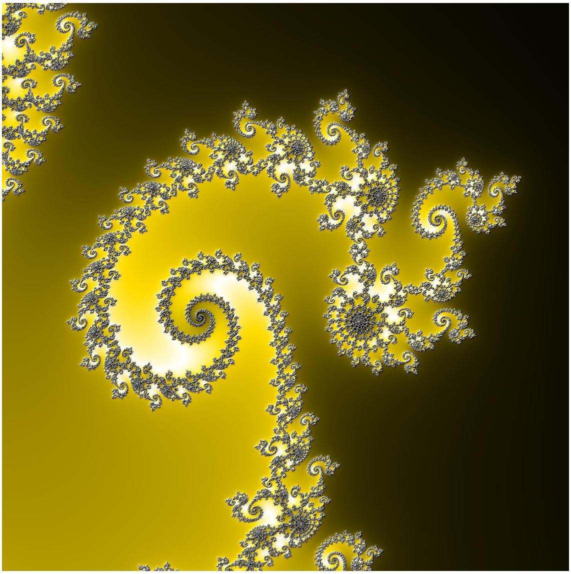

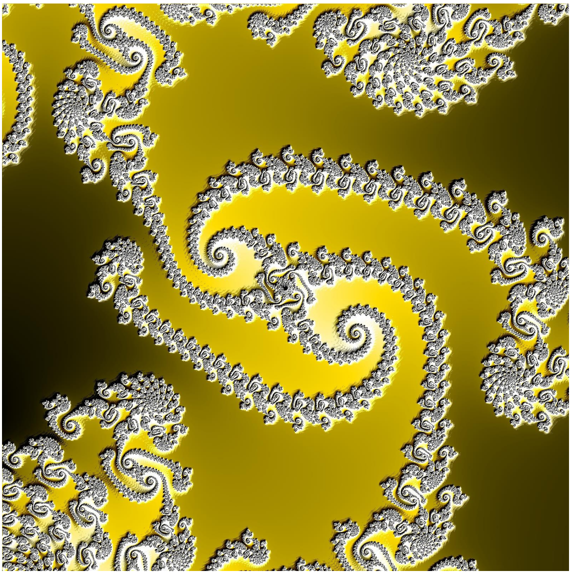

# Dendritic Structures

The image below is lobe 1:2:6 on the largest of the minibrots on the principal axis. You can tell that it is lobe 6 because of the principal synapse of order 6 and you can tell that it is on lobe 1:2:6 because of the order 2 spiral synapse. It is the long tendrils which tell you that this is not on the main brot but on one of the axonal minibrots.

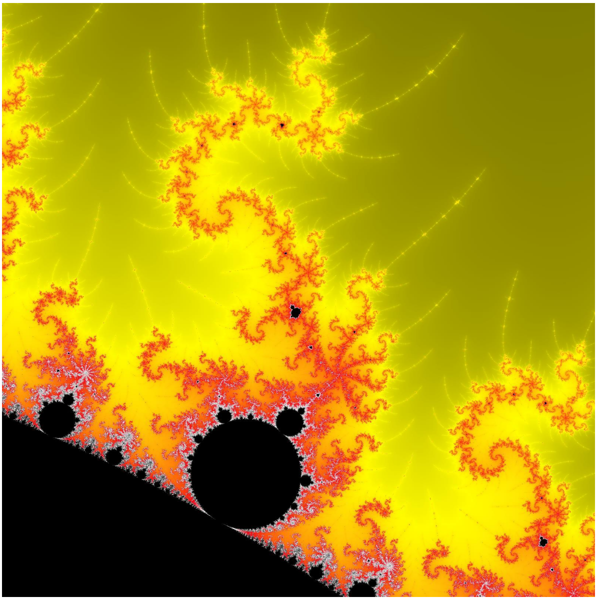

and here is a diamond brooch in one of the tendrils.

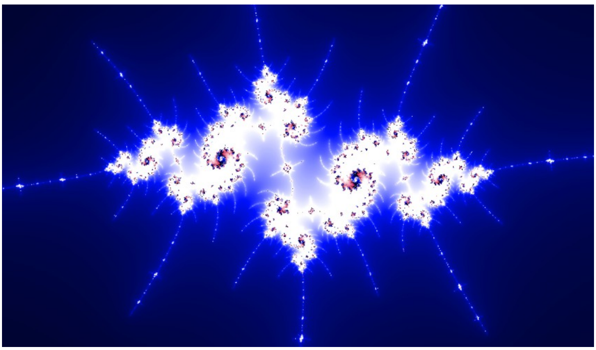

at whose centre we may find another brooch

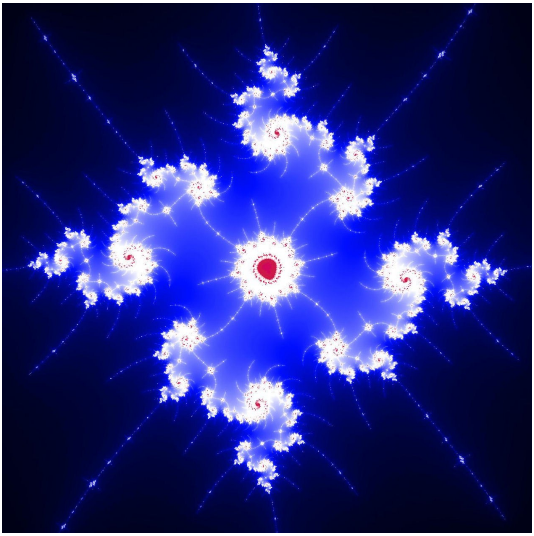

and a lovely ring!

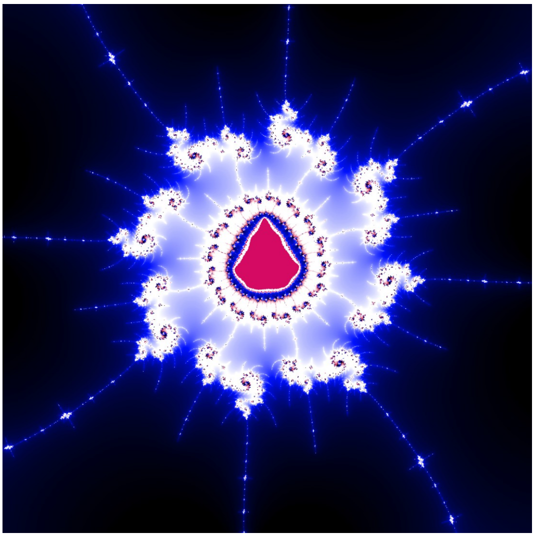

Moreover, there are other gems waiting to be discovered elsewhere inside this structure.

For example, here is a beautiful spiral galaxy...

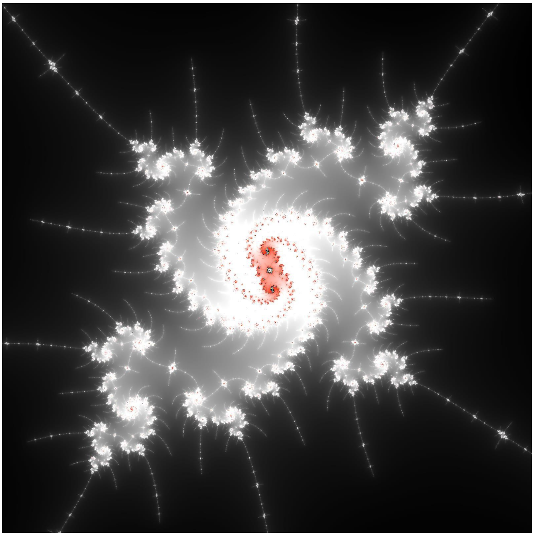

and here is an alien octopoid which I found at random.

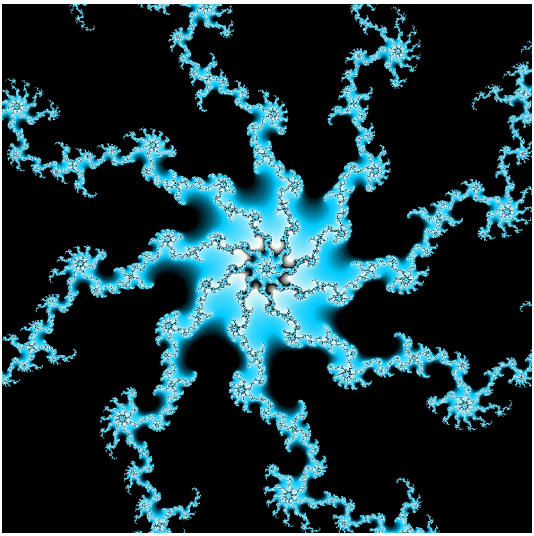

The magnification of this image was 53 trillion ($5.3 \times 10^{13}$). On this scale the whole map would be bigger than our galaxy!

# Axonal Structures

Because of the fact that axons consist only of synapses of order 2, structures on or near the axon of any neuron have a special symmetry. Let's have a look at the structures along the axon of the antenna.

It is difficult to see where the principal synapse is because all the synapses have order 2 and are perfectly straight but I believe it to be where I have drawn the arrow in the diagram on page 44.

All along the axon there are side 'sprouts' and at every point where the sprouts cross the axon there is a minibrot. It is easy to verify that these sprouts are the neurons of lobes 1:3 and 1:2›1 on each side.

The next most prominent sprouts emerge at 45 degrees and come from lobes 1:4 and 1:2:3 (not 1:3›2 as you might expect).

As you look at smaller and smaller minibrots, the surroundings of the minibrot become more and more symmetrical with rotational symmetry based on a binary division of the circle (i.e. 2, 4, 8 etc.).

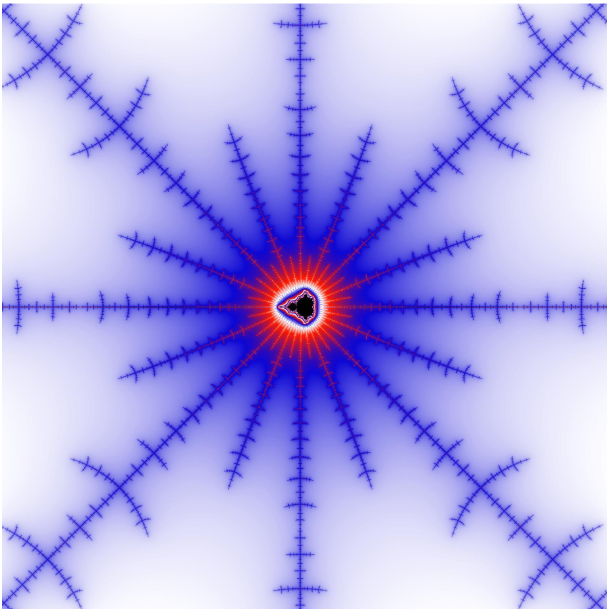

A particularly interesting series of cauliflower structures can be found as you move along the axis towards the cleft of the largest minibrot and some lovely  patterns can be found within the petals, all containing minibrots of their own.

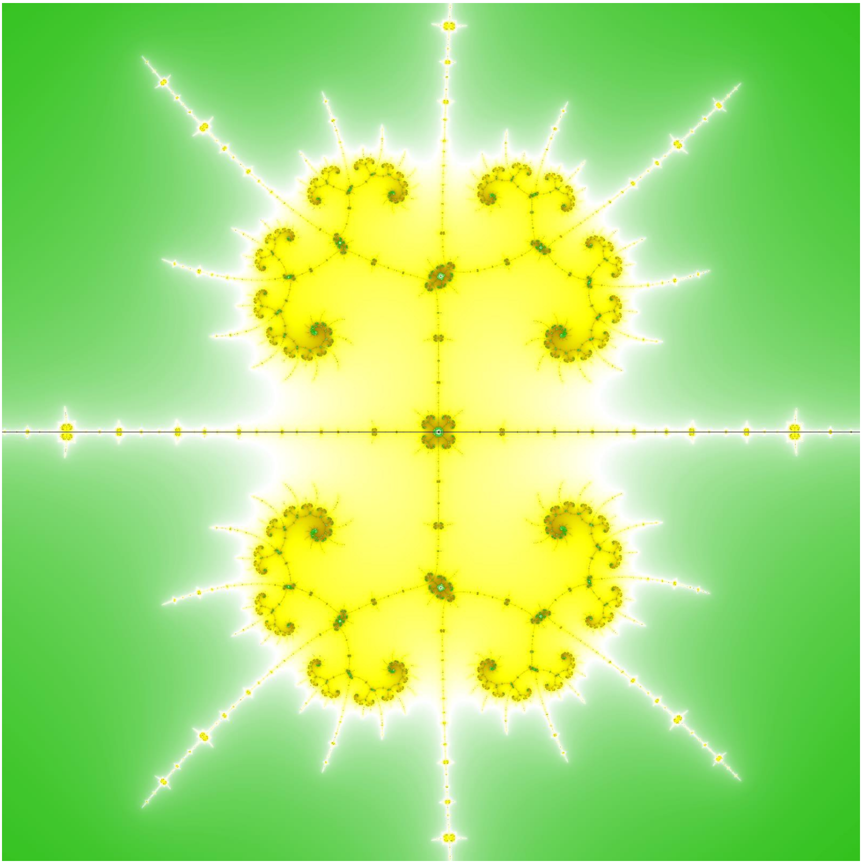

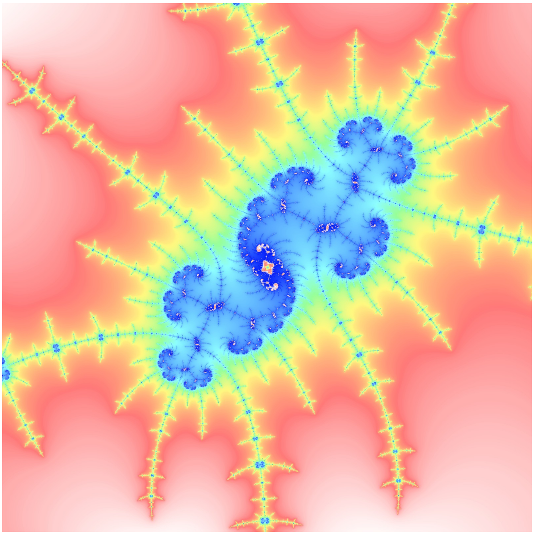

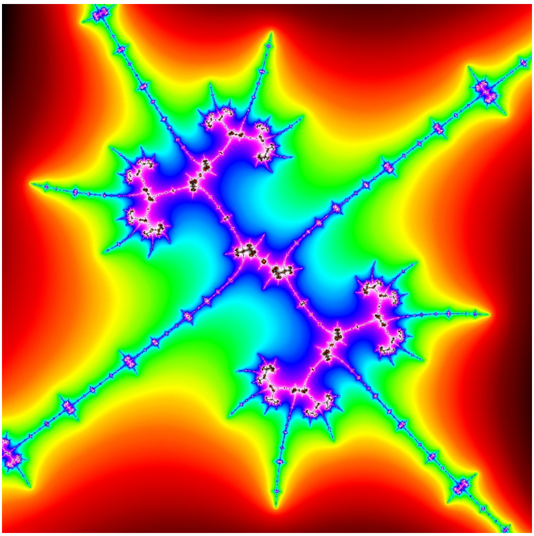

# Asymmetrical Features

While it is natural to want to zoom in on those features like minibrots and spirals which have a high degree of symmetry, it is worth seeking out places with other attractions. There is, for example, the procession of circus elephants which can be found in 'elephant valley' – the cusp of the cardioid:

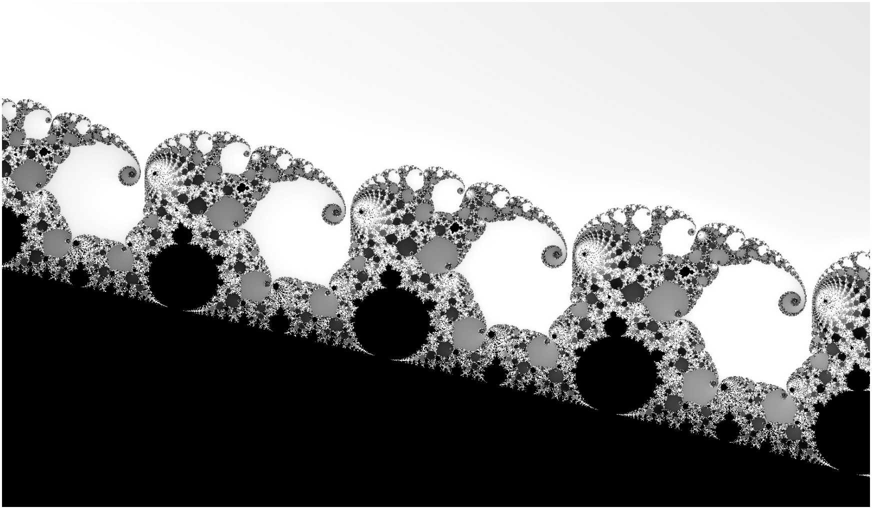

or the butterfly wings which can be found on the edge of lobe 4:

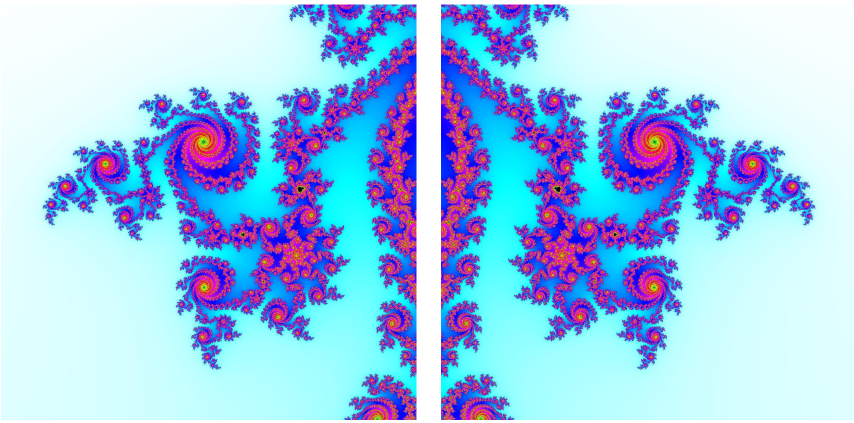

Another favourite image of mine is the dragons head on lobe 1:2:9:

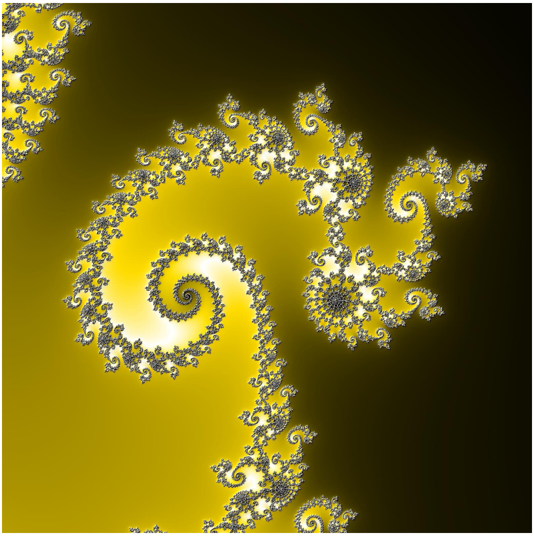

or is it a baroque wall bracket for a candle?

And what about the fire-breathing dragons of lobe 1:2:2?

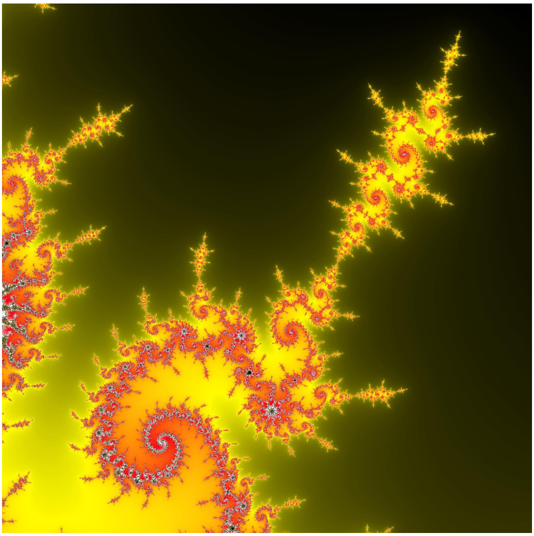

# Delving Deeper

I have already noted that, unlike Julia sets, the Mandelbrot map is not self-similar. The deeper you go you will find new structures appearing. It is, however, not easy to know where to start delving in search of new structures. It is all too tempting to zoom into a spiral synapse, for example, hoping to find something new only to find that it goes on for ever and ever. Or you might zoom into a minibrot only to be disappointed when you only find elephants and sea horse there. Mind you, the elephants and seahorses are not quite the same. These are to be found in the principal minibrot of lobe 3:

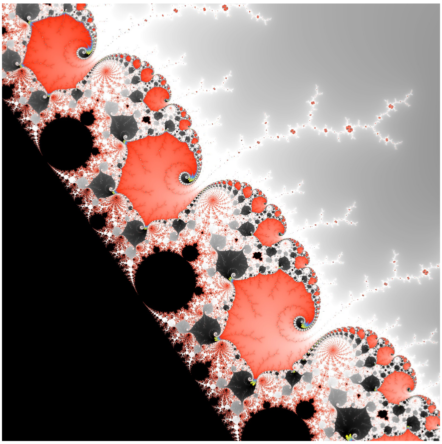

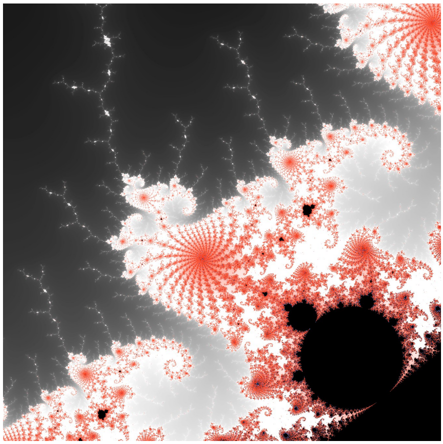

You will notice that they have both sprouted 'hair' with characteristic order 3 synapses. In addition, the 'hair' is decorated with beautiful brooches similar to but not the same as the ones that we found on the minibrots on the main axon.

If you choose any minibrot in one of the seahorses in seahorse valley (eg lobe 3›2$^3$) and then examine its elephants you will discover something like this:

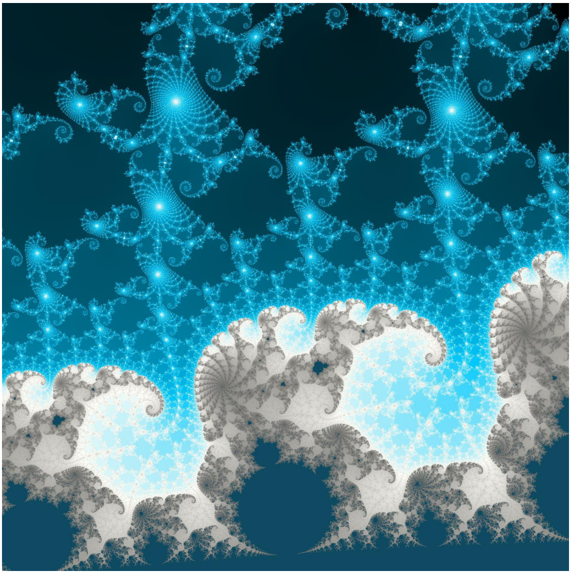

The elephants are surrounded by a swarm of seahorses!

I will leave you with the challenge of finding and naming the structure below:

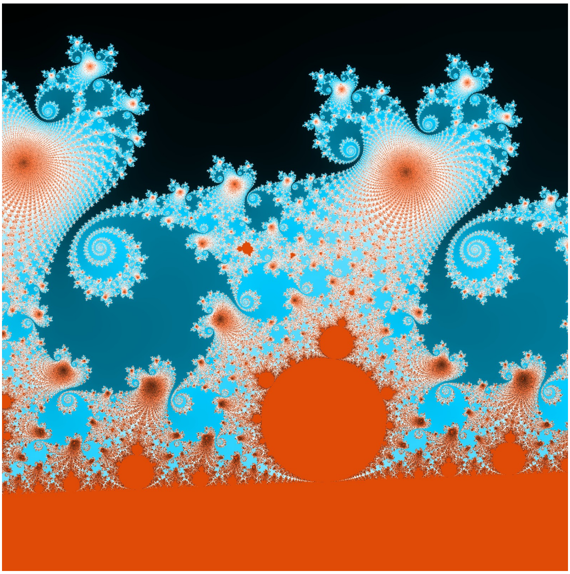f

# APPENDIX

*The Stern-Brocot tree theorem.*

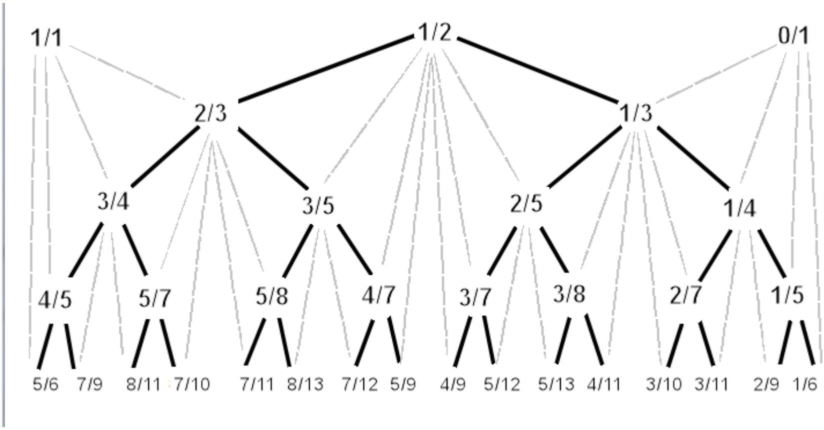


Every fraction $p/q$ in the Stern-Brocot tree has two parent fractions $a/b$ and $c/d$. Look closely and you will see that there is a simple relation between $a$, $b$, $c$ and $d$ namely

$$ad - bc = \pm 1 \tag{5}$$

It is easy to prove that this is always the case. Suppose that

$$\frac{a}{b} \text{¤} \frac{c}{d} = \frac{p}{q} \tag{6}$$

where $$p = a + c \quad \text{and} \quad q = b + d \tag{7}$$

We shall now show that the same relation holds between $a/b$ and $p/q$.

$$\begin{aligned} aq - bp &= a(b + d) - b(a + c) \\ &= ab + ad - ab - bc = 1 \end{aligned} \tag{8}$$

What this means is that for any pair of fractions for which equation (5) holds, the same relation will hold for all of its children.

Now since the relation holds for the two starting fractions 1/1 and 0/1, it must hold for all pairs of parental fractions in the tree.

Suppose we wish to find the two parents $a/b$ and $c/d$ for any fraction

$p/q$ where $p$ and q are co-prime.

We know from (5) that

$$aq - bp = 1 \tag{9}$$

hence

$$a = \frac{1 + bp}{q} \tag{10}$$

Now $(1 + bp)$ must be a whole multiple of $q$ and this multiple $a$ must be less that $p$ (because $a + c = p$). For example, if $p/q = 5/13$, then we are looking for a number in the sequence 6, 11, 16, 21 etc. which is a multiple of 13. In fact the next number in the sequence, the 5$^{th}$, is the one we are looking for – 26. This means that $a = 2$ and $b = 5$ from which it is easy to show that $c = 3$ and $d = 8$.

OK so the procedure worked for 5/13 but will it work for every fraction?

The problem boils down to the famous stepping stone problem. Suppose we put $q$ stepping stones in a circle (13 in this case) labelled 0 to $q - 1$. Starting from stone number 1, repeatedly jump $p$ steps at a time (in this case 5). Will you ever reach stone number 0 and if so, How many jumps will it take?

Well if $p$ and $q$ are co-prime it is easy to see that you will visit every stone once and once only in a unique cycle before returning to your starting point. What is more, you are bound to reach stone number 0 in less than $q$ jumps. In other words, the procedure will always find number $b$ ($< q$) and $a$ ($< p$) which will satisfy equation (9)

We have now proved that for every fraction $p/q$ there exists a parent $a/b$ where $a$ and $b$ are less than $p$ and $q$. But $a/b$ must have a parent too – and it must have a parent as well. Eventually we must reach either 1/1 or 0/1. And this essentially implies that every fraction must appear in the tree somewhere!

*Proof that lobe 2 is circular*

As stated on page 51 we require that

$$z^4 + 2cz^2 + c^2 + c = z \tag{11}$$

and

$$|4z^3 + 4cz| = 1 \tag{12}$$

Equation (11) is the general requirement for all points $z$ which have a cycle of period 2 for any value of $c$. Equation (12) is the requirement that the gradient of the $f$ function should be ±1.

It is not easy to solve equation (11) in terms of $z$ so lets solve it in terms of $c$. Fortunately, this simplifies quite a lot and we find that

either

$$c = -(z^2 - z) \tag{13}$$

or

$$c = -(z^2 + z + 1) \tag{14}$$

It turns out that equation (13) refers to the solutions appropriate to lobe 1. It is equation (14) which refers to lobe 2.

Solving this equation for $z$ gives us

$$z = \left(-1 \pm\sqrt{1 - 4(1 + c)}\right)/2 \tag{15}$$

which we shall write for convenience as

$$z = \left(-1 \pm\sqrt{X}\right)/2 \tag{16}$$

Now equation (12) can be rewritten as

$$|4z(z^2 + c)| = 1 \tag{17}$$

and using equation (14) we can eliminate $c$ to get

$$|4z(z + 1)| = 1 \tag{18}$$

Substituting (15) into (17) we get

$$\left|4\left(\frac{-1 \pm\sqrt{X}}{2}\right)\left(\frac{-1 \pm\sqrt{X}}{2} + 1\right)\right| = 1 \tag{19}$$

which eventually leads to

$$|(1 + c)| = 0.25 \tag{20}$$

Now if the modulus of 1 + $c$ is 0.25 then the locus of $c$ must be a circle of radius 0.25 centred on the point (-1, 0).

*Most of the illustrations in this book were generated using programs written by the author, many of which are available on his website: www.jolinton.co.uk.*