The MLC Conjecture in Complex Dynamics

Elliott Cawtheray
University of St Andrews
23rd April 2021

Abstract

This project focuses on one of the most important conjectures in complex dynamics - that the famous Mandelbrot set is a locally connected space. The conjecture is important because it would give huge amounts of information about the dynamics of all complex quadratic polynomials. For example, a proof of the conjecture solve a number of other problems, such as *density of hyperbolicity* for complex quadratics. This work first endeavours to build up the necessary background to state the conjecture, and then develops some of the basic theory. After this, we describe the cases where the conjecture has been settled, and delve into the mathematics used to get there. The final section explores a way that local connectivity of the Mandelbrot set can be leveraged, in particular to construct a homeomorphic abstract topological space known as the *pinched disc model.*

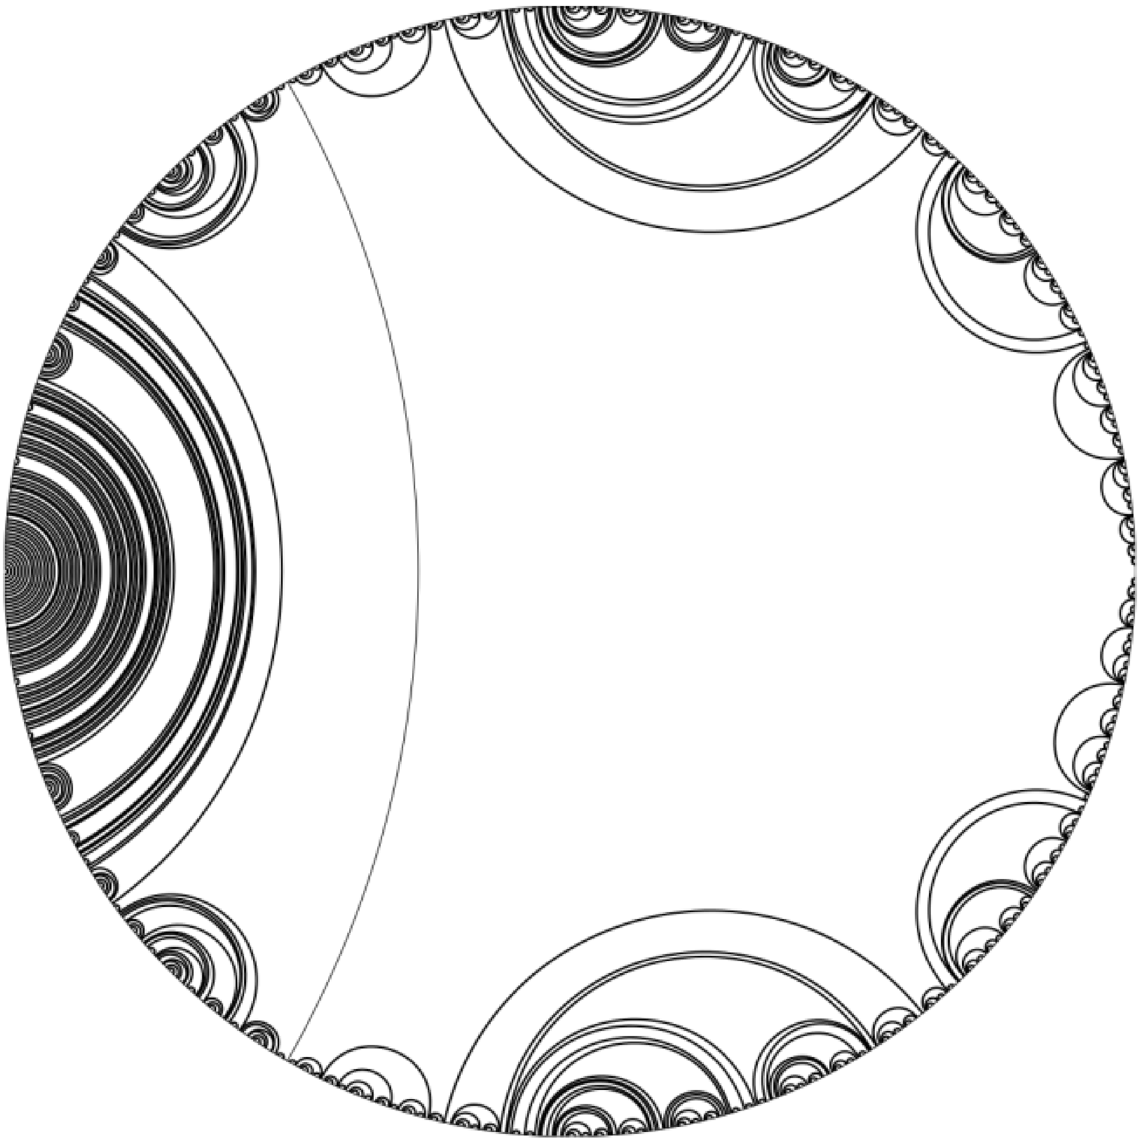

Thurston's Quadratic Minor Lamination [Thurston, 1985], which gives rise to a quotient space of the unit disc that is homeomorphic to the Mandelbrot set if and only if the Mandelbrot set is locally connected. This image was taken from [Blokh et al., 2017], where it is Figure 1 on page 3.

I certify that this project report has been written by me, is a record of work carried out by me, and is essentially different from work undertaken for any other purpose or assessment.

–Elliott Cawtheray

Contents

1 Introduction

"It can be argued that the mathematics behind these images is even prettier than the picture itself"
—Robert L. Devaney

The Mandelbrot set is one of the shining stars of mathematics. It fascinates with its beauty and its vast complexity, all resulting from a fairly simple process. It is worthy of the considerable attention and study that it has been the focus of, but in spite of this, some fundamental questions about it remain unanswered. The MLC conjecture, which states that the set is locally connected, is the most significant such unanswered question.

The set, which describes the dynamical behaviour of quadratic maps on the complex plane, had its properties studied and its importance outlined in the seminal works of Douady and Hubbard in the 1980s [Douady and Hubbard, 1985a], in which it was demonstrated (among many other things) that if the Mandelbrot set is locally connected, then the Carathéodory-Torhorst Theorem gives a parametrization of its boundary. Assuming local connectivity, this can be used to prove the conjecture that *hyperbolic maps* are dense in the space of complex quadratics. This *Density of Hyperbolicity Conjecture* (DHC)[1] has been worked on for over 100 years, since Fatou [Fatou, 1919]. The MLC conjecture's would-be consequences do not end here - if we can find a proof, various other conjectures about the dynamics of the quadratics would immediately follow.

Local connectivity has been shown at some points of the set. In 1989, Yoccoz proved local connectivity in the case of *non-infinitely renormalizable maps* [Hubbard, 1993], and much progress has been made for the remaining cases, most recently for some parameters of growing satellite type [Cheraghi and Shishikura, 2015], and parameters of stationary satellite type with specific combinatorics [Dudko and Lyubich, 2019]. At this level, the topic is immensely technical, but we make an attempt to explain what we can.

While the MLC conjecture holds the history of many decades of mathematics, the techniques that have been developed to tackle it have provided insights into a variety of families of maps, some of which are for the most part still to be studied. Thus, it is a conjecture with relevance to the future of mathematics as much as to the past.

This project was written with a certain audience in mind. A keen undergraduate who has good familiarity with university mathematics, and who has possibly even taken a class in topology, will derive maximal (if not necessarily non-zero) pleasure and use from this work. The opening sections define the Mandelbrot set and Julia sets, and investigate their basic properties, before moving on to connectedness and local connectedness, with plentiful examples. After this, we develop some of the theory of complex dynamics, including some deep results of complex analysis, and apply it to the family of functions that relate to the Mandelbrot set, leading to definitions of equipotentials and external rays. The project gets more advanced in the later sections, which explore the progress relating to the conjecture over the past few decades, as discussed above. The final section demonstrates how local connectedness of the Mandelbrot set can be applied, constructing the *pinched disc model* and using it to show that MLC implies DHC.

This work features no original results. The exposition is my own, though without exception based heavily on source material, which is cited at the start of the relevant sections. There are some original worked examples, and some original images which for the most part have been generated using the Python code linked below (making generous use of the in-built **matplotlib** and **PIL** modules - feel free to use/modify it!), and edited using the **GNU Image Manipulation Program**.

https://github.com/elliottmaths/MLC-Conjecture

[1]DHC for real quadratics has been proven by Graczyk and Świątek [Graczyk and Świątek, 1997], [Graczyk and Świątek, 1998] and Lyubich [Lyubich, 1997] (independently). In 2007, Kozlovski, Shen, and Strien [Kozlovski et al., 2007] proved DHC for the space of all real polynomials. The complex case, and even the complex quadratic case, remains open.

2 Background

2.1 Some housework

For most of this project, we will be looking at the complex plane. On those occasions where we are considering the real line, we will use the following familiar notation.

$$[a, b] = \{x \in \mathbb{R} : a \leq x \leq b\}, \quad (a, b) = \{x \in \mathbb{R} : a < x < b\}$$

To avoid confusion with the above, we will use angled brackets when working with ordered pairs, e.g. $\langle x, y\rangle \in X \times Y$ for $x \in X$, $y \in Y$.

Let $\mathbb{S}$ be a set of sets. Then $\bigcup \mathbb{S}$ denotes the union of the sets in $\mathbb{S}$. That is, $\bigcup \mathbb{S} = \{x \in X : X \in \mathbb{S}\}$.

Let $r \in \mathbb{R}^+$. We will use the following notation for the open ball on the complex plane, and its variants:

$$B(z_0, r) = \{z \in \mathbb{C} : |z - z_0| < r\}, \quad C(z_0, r) = \{z \in \mathbb{C} : |z - z_0| = r\}, \quad D(z_0, r) = \{z \in \mathbb{C} : |z - z_0| \leq r\}$$

which can be handily remembered as the open **b**all, the **c**ircle, and the close**d** ball.

Unless otherwise stated, we will be working with the standard topology on the complex plane induced by the metric $d(z, w) = |z - w|$, where $|\cdot|$ is the complex modulus. For example, for any $A \subseteq \mathbb{C}$, we define the diameter of A to be $\operatorname{diam} A = \{\sup |z - w| : z, w \in A\}$, the largest possible distance between any pair of points in A.

We take $\mathbb{N} = \{0, 1, 2, \ldots\}$, and will use $\mathbb{Z}^+$ when we wish to exclude 0.

Derivatives will be denoted with a dash, e.g. f'. If $f : A \to B$ is a function, and $\mathcal{A} \subset A$, we denote the restriction of f to $\mathcal{A}$ by $f|_{\mathcal{A}}$, which is the function $\mathcal{A} \to B$ given by $f|_{\mathcal{A}}(x) = x$ for all $x \in \mathcal{A}$.

Let X be a topological space. We denote its closure by $\overline{X}$, and its interior by $\operatorname{int}(X)$. If we have an equivalence relation $\sim$ defined on X, we write $[x]$ for the subset of X consisting of all elements of X that are equivalent to a given $x \in X$, and call this subset the equivalence class of X containing x. This partitions the elements of X, and we can obtain the quotient space $X/\!\sim\, = \{[x] : x \in X\}$ which is given the finest possible topology such that the map $x \mapsto [x]$ is continuous.

We will often be looking at the parametrization $e^{2\pi i t}$ of $C(0, 1)$, where t runs from 0 to 1. We will say t is a member of $\mathbb{R}/\mathbb{Z}$, which is the unit interval with 0 and 1 identified. Note that distances in $\mathbb{R}/\mathbb{Z}$ are given by $d_{\mathbb{R}/\mathbb{Z}}(e^{2\pi i s}, e^{2\pi i t}) = \min\{|t - s|, |1 - t - s|\}$, which is the minimum distance between the two points in $C(0, 1)$. When we want to look at this parametrization for rational t, we will say $t \in \mathbb{Q}/\mathbb{Z}$, which is the set of rational values between 0 and 1, with 0 and 1 identified. If $p/q \in \mathbb{Q}/\mathbb{Z}$, we will always take p and q to be co-prime.

When otherwise not stated, take c to be an arbitrary point in the parameter plane throughout this work.

2.2 The Mandelbrot set

"In the whole of science, the whole of mathematics smoothness was everything. What I did was to open up roughness for investigation."
—Benoit B. Mandelbrot

Due to its intricate structure, there are multiple ways one could think of what the Mandelbrot set actually is. We will be getting in to many of these, but to define it, we will start by considering it as a classification of complex quadratic functions based on their dynamical behaviour. To this end, it will be useful to introduce something called the parameter plane, which is where the Mandelbrot set will live. It is the set $\mathbb{C}$ of complex numbers, where we will think of each point c as a parameter, which will be used to define a function. In particular, for each $c \in \mathbb{C}$, define the function $f_c : \mathbb{C} \to \mathbb{C}$ by $f_c(z) = z^2 + c$, so that we have the family of complex functions

$$F = \{f_c : c \in \mathbb{C}\}$$

This family is of particular importance because it is able to describe the dynamical behaviour of all complex quadratics, while also consisting of elements of a 'nice' form. F is sometimes called the Douady-Hubbard family of quadratic polynomials after Adrien Douady and John Hamal Hubbard whose foundational work [Douady and Hubbard, 1985a] studied the dynamics of elements of F and proposed the conjecture on which this work focuses. Note that all $f_c \in F$ are holomorphic.

So, we have the parameter plane, where each point of the complex plane has a quadratic function associated with it. Let us briefly investigate the dynamical behaviour of the elements of F, that is, what happens to

$$f_c^{\circ k}(z) := (\underbrace{f_c \circ ... \circ f_c}_{\text{k times}})(z)$$

for some $c, z \in \mathbb{C}$, and as $k \to \infty$. The fact that c and z both come from a complex plane is a potential cause for confusion, and it is for this reason we introduced the idea of a parameter plane, which c varies over. Analogously, we say the complex plane over which z varies is called the dynamical plane. So, to reiterate[2], we have the parameter plane where each point of the complex plane is associated with a function f_c defined by the parameter c, and then for any given c, the domain of f_c is also a complex plane that we call the dynamical plane. Choosing a starting point $z_0 \in \mathbb{C}$ in the dynamical plane, for a fixed $c \in \mathbb{C}$ in the parameter plane, the point z_0 will jump around the dynamical plane by the sequence $z_0, f_c(z_0), f_c^{\circ 2}(z_0), ...$, and we have dynamics, explaining our choice of terminology.

Example 2.1 Take $c = 0$, so that we are considering the function $f_0 : z \mapsto z^2$. The general dynamical behaviour is easily described. If $|z_0| < 1$, then the iterates $|f_0^{\circ k}(z_0)| \to 0$ as $k \to \infty$. If $|z_0| > 1$, then $|f_0^{\circ k}(z_0)| \to \infty$. Finally, if $|z_0| = 1$, then $|f_0^{\circ k}(z_0)| = 1$ for all $k \in \mathbb{N}$, so that our iterates jump around the unit circle in the plane (this case is where the most interesting dynamics arise for this quadratic!).

Example 2.2 Suppose $c \in \mathbb{C}$ satisfies $(c^2 + c)^2 + c = 0$. One of the solutions to this quartic is $\omega \approx -0.12 + 0.745i$. Then, the sequence of iterates will go $0, \omega, \omega^2 + \omega, (\omega^2 + \omega)^2 + \omega = 0, \omega, \ldots$, so that we have a cycle of length 3. In fact, even if we approximate our parameter and perturb our initial point slightly away from 0, the sequence of iterates appears to be pulled towards this cycle (see Fig 1). We will return to this example regularly, and for the duration of this paper, ω is fixed to be this parameter.

[2] No pun intended.

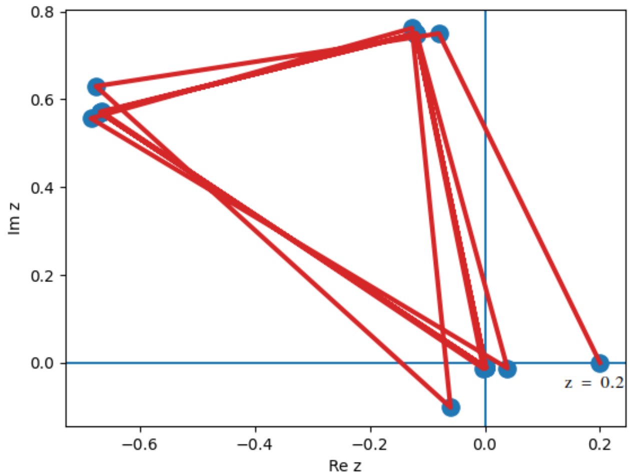

Figure 1: The first 30 iterates of f_ω on the function's dynamical plane, starting at the point $z = 0.2$.

Example 2.3 Take $c = -0.7 - 0.3i$, with our starting point $z_0 = 0$. At first, the iterates seem to calmly jump around inside the circle $|z| = 2$. The 25^{th} iterate escapes this circle, and from there the iterates blow up - In Figure 2, we show only the iterate immediately following this as otherwise it's too big to scale. The core takeaway is that $|f_c^{\circ k}(0)| \to \infty$ as $k \to \infty$.

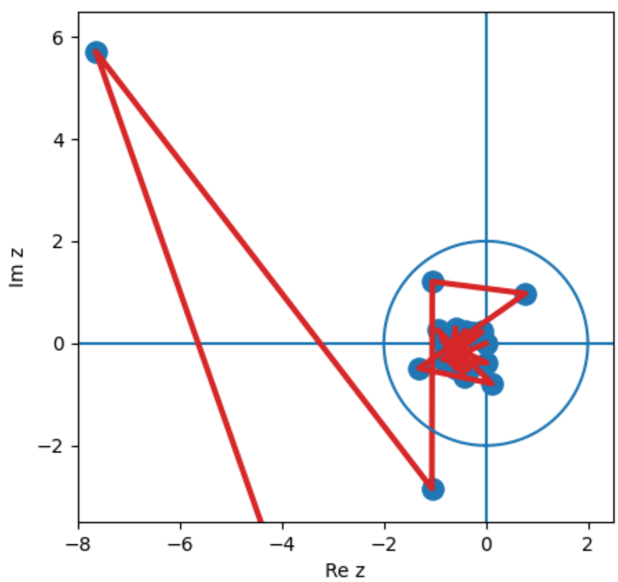

Figure 2: The first 26 iterates of f_c on the function's dynamical plane, with $c = -0.7 - 0.3i$ and starting at the point $z = 0$. The circle $C(0, 2)$ is also shown, with the iterates remaining tame and inside it to begin with. Once an iterate escapes the circle, the subsequent iterates explode in magnitude.

At the start of this section, it was claimed that the family F fully describes the dynamics of all complex quadratic maps. We formalise this a little bit. Suppose we have two functions $f, g : \mathbb{C} \to \mathbb{C}$ such that there exists a function $h : \mathbb{C} \to \mathbb{C}$, with $f = h^{-1} \circ g \circ h$. Then

$$f^{\circ k} = (h^{-1} \circ g \circ h)^{\circ k} = (h^{-1} \circ g \circ h) \circ \ldots \circ (h^{-1} \circ g \circ h) = h^{-1} \circ g^{\circ k} \circ h$$

so that the dynamics of f and g are in some way linked. We call h a change of coordinates, and say f and g are conjugate to each other. Depending on the context, we may want h to satisfy some property, so that we have a good handle on the way that the dynamics of our two maps are linked. In our case, we will have $f : z \mapsto a_2 z^2 + a_1 z_1 + a_0$ as an arbitrary complex quadratic, $g = f_c$ for some $f_c \in F$

(uniquely determined by f), and h as an affine change of coordinates, so that $h : z \mapsto az + b$ for some $a, b \in \mathbb{C}$. Then, for the purposes of our topological and dynamical considerations, the maps f and f_c are effectively equivalent. In particular, it is clear that with such an h, iterates of f stay bounded if and only if iterates of f_c do. It remains to find our h. Suppose $f = h^{-1} \circ f_c \circ h$. Then

$$\begin{aligned} f(z) = a_2 z^2 + a_1 z + a_0 &= (h^{-1} \circ f_c \circ h)(z) = (h^{-1} \circ f_c)(az + b) = h^{-1}((az+b)^2 + c) \\ &= h^{-1}(a^2 z^2 + 2abz + b^2 + c) = \frac{a^2 z^2 + 2abz + b^2 + c - b}{a} = az^2 + 2bz + \frac{b^2 - b + c}{a} \end{aligned}$$

for any $z \in \mathbb{C}$, so that $a_2 = a$, $a_1 = 2b$, $a_0 = \frac{b^2 - b + c}{a}$, and we can uniquely determine $c = a_2 a_0 - \frac{a_1^2}{4} + \frac{a_1}{2}$.

It was also suggested that F's elements were in some way nice to work with. There are obvious benefits to working with only one parameter, but it goes further than this. A critical point of a function $f : \mathbb{C} \to \mathbb{C}$ is a point at which $f'(z) = 0$. We shall see later that taking critical points as initial points and looking at the behaviour of the resulting iterates can provide us with a lot of information. For now, we note that all $f_c \in F$ have a single critical point, at $z = 0$. The image of a critical point is called a critical value. Thus, each $f_c \in F$ has the single critical value c.

If we are willing to take on trust for now that $z = 0$ is in some way a special point to look at in the dynamical plane, our definition of the Mandelbrot set is a logical one:

Definition 2.4 $M = \{c \in \mathbb{C} : |f_c^{\circ k}(0)| \not\to \infty \text{ as } k \to \infty\}$.

That is, M consists of all the points $c \in \mathbb{C}$ in the parameter plane whose associated function f_c does not go off to infinity as we iterate more and more, starting at the critical point $z = 0$ of f_c.

The set was named after fractal geometry legend Benoit B. Mandelbrot[3] by Douady and Hubbard [Douady and Hubbard, 1985a].

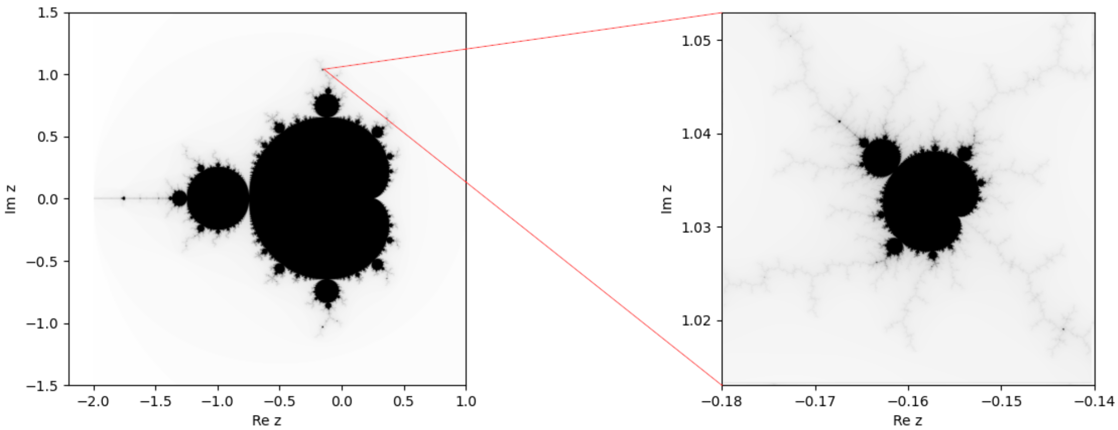

Figure 3: Left: The Mandelbrot set is pictured on the parameter plane. More accurately, the Mandelbrot set is being approximated. We turn some amount of pixels into points in the parameter plane, and test how many iterates it takes for them to leave $D(0, 2)$ (see Corollary 2.6). The more iterates, the darker the pixel.

Right: We zoom in on the point $c \approx -0.16 + 1.035i$, and find an approximate copy of the Mandelbrot set within. The presence of these 'baby Mandelbrots' within itself will be explained in Section 4.3.

Figure 3 is one of the most famous images in all of mathematics, and for good reason - it is stunning and infinitely intricate, in spite of its simple definition. We briefly establish a couple of M's properties.

[3]The 'B.' stands for Benoit B. Mandelbrot

Theorem 2.5 *The Mandelbrot set is bounded, in particular $M \subseteq D(0,2)$.*

Proof. We begin with a **Lemma**: For $|c| > 2$, the sequence of (absolute values of) iterates $\{\left|f_c^{\circ k}(0)\right|\}_{k\in\mathbb{N}}$ is monotonically increasing. The proof, which we omit, is an exercise in strong induction. Next, suppose $|c| > 2$. Then by our lemma, we may define $\varepsilon_k = \left|f_c^{\circ k}(0)\right| - 2 > 0$ for each $k \geq 1$, and will have $\varepsilon_1 < \varepsilon_2 < \varepsilon_3 < \dots$

Therefore,

$$\begin{aligned}
\left|f_c^{\circ k}(0)\right| &= \left|f_c(f_c^{\circ(k-1)}(0))\right| = \left|f_c^{\circ(k-1)}(0)^2 + c\right| \geq \left|f_c^{\circ(k-1)}(0)\right|^2 - |c| = \left|f_c^{\circ(k-1)}(0)\right|^2 - |f_c(0)| \\
&\geq \left|f_c^{\circ(k-1)}(0)\right|^2 - \left|f_c^{\circ(k-1)}(0)\right| = \left|f_c^{\circ(k-1)}(0)\right| \left(\left|f_c^{\circ(k-1)}(0)\right| - 1\right) = \left|f_c^{\circ(k-1)}(0)\right| (\varepsilon_{k-1}+1) \\
&\geq \left|f_c^{\circ(k-2)}(0)\right| (\varepsilon_{k-1}+1)(\varepsilon_{k-2}+1) \geq \dots \geq |f_c(0)| \prod_{r=1}^{k-1} (\epsilon_r + 1) \\
&\geq |f_c(0)| \left(\varepsilon_1 + 1\right)^{k-1} \to \infty \text{ as } k \to \infty
\end{aligned}$$

so that $\left|f_c^{\circ k}(0)\right| \to \infty$ as $k \to \infty$, and $c \notin M$. □

Corollary 2.6 *If $|f_c^{\circ k}(0)| > 2$ for any $k \in \mathbb{N}$, then $c \notin M$. Therefore,*

$$M = \{c \in \mathbb{C} : |f_c^{\circ k}(0)| \leq 2 \textit{ for all } k \in \mathbb{N}\}.$$

Theorem 2.7 *The Mandelbrot set is compact, and has a non-empty interior.*

Proof. For compactness, it suffices to show that M is closed, by the Heine-Borel Theorem adapted to the complex plane, and by Theorem 2.5. For each $n \in \mathbb{Z}^+$, define

$$A_n = \{c \in \mathbb{C} : |f_c^{\circ n}(0)| \leq 2\} = (f_c^{\circ n})^{-1}\left(D\left(0,2\right)\right)$$

As the pre-image of a closed set under a continuous function, A_n is closed. Further, if we use Corollary 2.6, we get that

$$\bigcap_{n\in\mathbb{Z}^+} A_n = \{c \in \mathbb{C} : |f_c^{\circ n}(0)| \leq 2 \text{ for all } n \in \mathbb{Z}^+\} = M$$

As an intersection of closed sets, M must be closed.

Looking to the interior, suppose $|c| \leq \frac{1}{4}$. Then

$$|f_c(0)| = |c| \leq \frac{1}{4} \leq \frac{1}{2}$$

Now suppose $\left|f_c^{\circ l}(0)\right| \leq \frac{1}{2}$ for all $l < k$. Then

$$\left|f_c^{\circ k}(0)\right| = \left|f_c^{\circ(k-1)}(0)^2 + c\right| \leq \left|f_c^{\circ(k-1)}(0)\right|^2 + |c| \leq \frac{1}{4} + \frac{1}{4} = \frac{1}{2}$$

Thus, $\left|f_c^{\circ k}(0)\right| \leq \frac{1}{2} \leq 2$ for all $k \in \mathbb{N}$, and by Corollary 2.6, $c \in M$. □

Looking back at Examples 2.1, 2.2, and 2.3, we see that $0 \in M$ and $-0.7-0.3i \notin M$, with Corollary 2.6 explaining the presence of the circle $C(0,2)$ in Figure 2. Finally, $\omega \in M$, and in fact all solutions to $(c^2+c)^2 + c = 0$ belong to M, since the iterates stay bounded due to their cyclic nature.

It was conjectured by Mandelbrot himself in 1985 [Mandelbrot, 1985] that the boundary of the Mandelbrot set, which has a fractal nature, has Hausdorff dimension 2, and this result was later proved by Shishikura [Shishikura, 1998].

Julia sets

What if we are to start iterating from a point in the dynamical plane other than $z = 0$? In particular, which choices of initial point will iterate to infinity for a given $f_c \in F$, and which will stay bounded? We are motivated to define the *filled-in Julia set* of f_c below:

Definition 2.8 $K(c) = \{z \in \mathbb{C} : |f_c^{\circ k}(z)| \not\to \infty \text{ as } k \to \infty\}$
so-called because we will be especially interested in its boundary, the *Julia set* of f_c, again defined below:

Definition 2.9 $J(c) = \partial K(c)$.

Note that $K(c)$ and $J(c)$ live in the dynamical plane, in contrast to M. So, $K(c)$ is the set of initial points from which we may iterate f_c and stay bounded, while $J(c)$ consists of points in the plane that are arbitrarily close to initial points that stay bounded, and initial points that don't. We have another (fairly tautological...) characterisation of the Mandelbrot set as $M = \{c \in \mathbb{C} : 0 \in K(c)\}$.

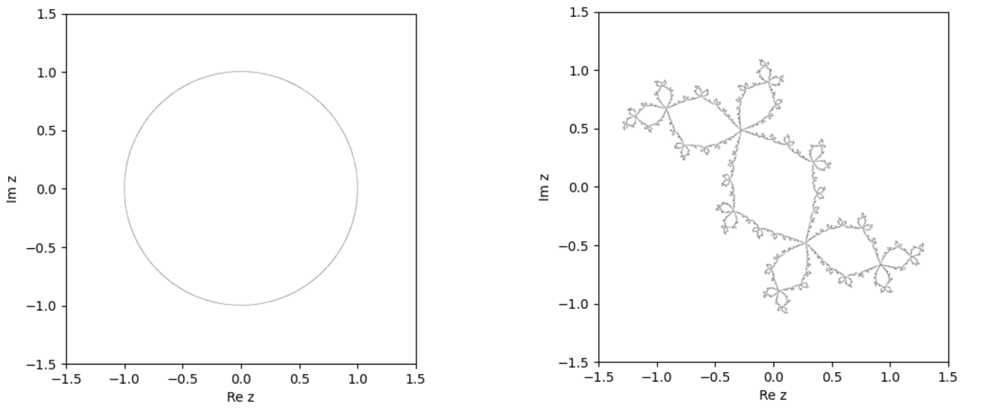

Figure 4: Left: The Julia set for $z \mapsto z^2$, as in Example 2.1. In this case, $J(0) = C(0, 1)$.

Right: The Julia set for $z \mapsto z^2 + \omega$, as in Example 2.2. This Julia set is known as the Douady rabbit.

It may be shown that $J(c)$ is non-empty and compact for any $c \in \mathbb{C}$. The following is of dynamical importance.

Proposition 2.10 *Let* $f_c \in F$*. Then* $J(c) = f_c(J(c)) = f_c^{-1}(J(c)) = \{z \in \mathbb{C} : f_c(z) \in J(c)\}$*. That is,* $J(c)$ *is completely invariant under* f_c*.*

Proof. $z \in \mathbb{C}$ stays bounded under iteration if and only if $f_c(z)$ stays bounded under iteration. Note that $f_c^{-1}(z)$ will consist of two complex numbers, unless $z = c$. In any case, the elements of $f_c^{-1}(z)$ stay bounded under iteration if and only if z does, since they both go to z. Therefore, $z \in K(c)$ if and only if $f_c(z) \in K(c)$, if and only if $f_c^{-1}(z) \subseteq K(c)$. This is equivalent to saying $K(c) = f_c(K(c)) = f_c^{-1}(K(c))$. Given the filled-in Julia set is not changed at all by forward or backwards iteration, the same must be true of its boundary $J(c)$. □

2.3 Connectivity and the Mandelbrot set

Notions of connectivity in both the dynamical and parameter planes are at the heart of the key questions in holomorphic dynamics, which is something you might not expect at first. We look in particular at the notion of a connected space, allowing for a discussion of the 'dichotomy' imposed on Julia sets for quadratic maps. Then we move on to local connectivity, for now only defining it and looking at a few examples. In later sections, we will see the huge ramifications of local connectedness in the parameter space.

Definition 2.11 A topological space X is *connected* if there is no pair of non-empty, disjoint open sets U, V such that $X = U \cup V$.

Example 2.12 Let $A \subseteq \mathbb{R}$. Then A is connected if and only if A is an interval.

It is a deep result of Douady and Hubbard that M is connected. We note it now. A proof may be found in Chapter 8 of [Douady and Hubbard, 1985a].

Theorem 2.13 *M is connected.*

The Julia sets pictured in Figure 4 are both connected. We note that both have corresponding parameters belonging to M.

A subset of a space X is connected if it is connected as a subspace. In any space, the singletons (subsets consisting of a single element) are always connected subsets, giving reflexivity for the following equivalence relation.

Definition 2.14 Let X be a topological space. Define an equivalence relation $\sim$ on X by $x \sim y$ if and only if there is a connected subspace of X that contains x and y. Thus, we partition the space into maximal connected subsets, that we call *connected components*.

So, a space X is connected if and only if $x \sim y$ for all $x, y \in X$. We might instead have the opposite situation.

Definition 2.15 A topological space X is *totally disconnected* if $x \sim y \Rightarrow x = y$. That is, if the connected components are the singletons, and no two distinct points belong to the same connected component.

In a totally disconnected space, the only connected subsets are the singletons and (vacuously) the empty set.

Example 2.16 (Cantor middle-third set) In the following, for a set A and an $x \in \mathbb{R}$, define $x + \frac{A}{3} = \{x + \frac{a}{3} : a \in A\}$. Let $C_0 = [0, 1]$, and for each $n \in \mathbb{N}$, define the set

$$C_n = \frac{C_{n-1}}{3} \bigcup \left(\frac{2}{3} + \frac{C_{n-1}}{3}\right)$$

Then the Cantor middle-third set is defined as

$$C = \bigcap_{n=0}^{\infty} C_n$$

We can think of this construction as taking $[0, 1]$, and removing its middle third. Then, we remove the middle thirds of the two resulting subsets. Repeating this process countably many times, we obtain C. It is a famous set in mathematics, due to the following properties: it is compact, totally disconnected and perfect (it is closed and has no isolated points). In fact, any set that satisfies these 3 properties is homeomorphic to C. As a result, we refer to any such set as a Cantor set, and all Cantor sets are topologically 'the same'.

Let $c < -2$. Then c is not in M, by Corollary 2.6. The Julia set $J(c)$ will be a Cantor set contained in the real line. The Julia set from Example 2.3, corresponding to the parameter $-0.7 - 0.3i \notin M$, is totally disconnected, and in fact a Cantor set. Click here for an animated visualisation of this Julia set.

The remarkable observation of Julia [Julia, 1918] and Fatou [Fatou, 1919] was that these examples are no coincidence. The following great theorem gives rise to an alternative definition of the Mandelbrot set which, on the face of it, is radically different to how we have been thinking thus far.

Theorem 2.17 *Let $f_c \in F$. Suppose $0 \in K(c)$, i.e. $f_c^{\circ k}(0) \not\to \infty$. Then $J(c)$ is connected. Suppose instead $0 \notin K(c)$, so that $f_c^{\circ k}(0) \to \infty$. Then $J(c)$ is not connected, and is in fact a Cantor set.*

This is often referred to as the Dichotomy theorem, as it gives rise to a clear dichotomy of Julia sets for the functions of F, where they are either composed of a single piece, or of uncountably many disjointed parts, and (incredibly) this is determined by the behaviour of the critical point under iteration by the function in question. Note that the connectivity of $J(c)$ and $K(c)$ are equivalent.

Corollary 2.18 $M = \{c \in \mathbb{C} : J(c) \text{ is connected}\}$.

As a result of this corollary, the Mandelbrot set is sometimes referred to as the connectedness locus of the quadratics.

Our last definition here is that of simple connectivity. The concept will be of no immediate use to us, but will become important later (see Section 4.3). We give a super informal definition, since to do otherwise would bring us firmly off-topic for no real benefit. Later on, in Corollary 4.7, we will have a more rigorous characterisation of simple connectivity, which will be entirely sufficient for our purposes.

Definition 2.19 A topological space X is *simply connected* if it is 'path connected', and if it 'contains no holes'. In the particular case of subsets of the complex plane, this means that we may connect any pair of points in the space by a curve whose image is contained in the space, and that any loop drawn in the space may be continuously shrinked to a point.

Example 2.20 The complex plane $\mathbb{C}$ is seen to be simply connected, as is the unit ball $B(0, 1)$.

Local connectivity

We are as far as the 'M' of 'MLC'. Now we turn to the 'LC', and investigate its nuances with some examples.

Definition 2.21 A topological space X is *locally connected* at $x \in X$ if, for every open set U containing x, there is a connected open set $V \subseteq U$ that still contains x. X is a locally connected space if it is locally connected at all $x \in X$.

Note that U and V are open sets *in the space* X. It is clear that we will have local connectedness at a point if there are arbitrarily small open neighbourhoods of that point which are connected.

Perhaps surprisingly, connectedness and local connectedness are logically independent of one another, in that whether or not a space is connected has no bearing on whether or not it is locally connected, and vice versa. We illustrate this through the following four examples.

Example 2.22 The set $[0, 1]$ is connected and locally connected. To see it is locally connected, let $x \in [0, 1]$, and let U be some open neighbourhood of x in $[0, 1]$. Then U open implies $[0, 1] \cap (x-\varepsilon, x+\varepsilon) \subseteq U$ for some $\varepsilon > 0$. But then this interval is a connected open set containing x, that is contained in U, so that $[0, 1]$ is locally connected at x. Since x was arbitrary, the space $[0, 1]$ is locally connected.

Example 2.23 The Cantor set C, as in Example 2.16, is neither connected nor locally connected. In fact, a totally disconnected space is locally connected if and only if it is discrete. For, given a point x in a totally disconnected space X, the only connected subset of X which contains x is $\{x\}$, so that points can have connected neighbourhoods at all if and only if all the singletons are open neighbourhoods, and this is equivalent to the space being discrete. Conversely, if the space is discrete, then these one-point neighbourhoods are clearly arbitrarily small open neighbourhoods.

Example 2.24 The set $(-1, 0) \cup (0, 1)$ is not connected, by construction, but is locally connected by the same argument as in Example 2.22.

We have to work a little bit harder to give a space that is connected but not locally connected, but it's well worth it.

Example 2.25 (Topologist's sine curve) Consider the continuous function $f : (0,1] \to \mathbb{R}$ given by $f(x) = \sin\frac{1}{x}$, and in particular its graph $G = \{\langle x, f(x)\rangle : x \in (0,1]\}$, considered as a subspace of $\mathbb{R}^2$. Since the graph of a continuous function is homeomorphic to its domain, G is homeomorphic to the interval $(0,1]$, and in particular is connected. Note also that the closure of G is

$$\overline{G} = G \cup (\{0\} \times [-1,1])$$

since for any $\langle 0,t\rangle \in \{0\} \times [-1,1]$, the sequence $\{\langle \frac{1}{\sin^{-1}(t)+2\pi n}, t\rangle\}_{n\in\mathbb{N}}$ in G tends to $\langle 0,t\rangle$. $\overline{G}$ is called the topologist's sine curve, and serves as an interesting example in topology.

We will first show that it is connected. Suppose otherwise, so that $\overline{G} = A \cup B$ for open, disjoint and non-empty $A, B \subseteq \overline{G}$. Note that $G \subseteq \overline{G}$ implies that $G = G \cap \overline{G} = (G \cap A) \cup (G \cap B)$. But G is connected, so that one of the sets in this union must be empty. Without loss of generality, suppose $G \cap A = \varnothing$. Now, A is non-empty, so that there is some $a \in A \subseteq \overline{G}$. Since $a \notin G$, it must be a limit point of G, and so any neighbourhood of a has non-trivial intersection with G. But a belongs to the open set A, so that there is an open neighbourhood U of a which is contained in A and therefore disjoint from G. We have a contradiction, and thus $\overline{G}$ is connected.[4]

However, $\overline{G}$ is not locally connected at any point $\langle 0,t\rangle$ for $t \in [0,1]$, and thus $\overline{G}$ is not a locally connected space. For example, consider a sufficiently small open ball about $\langle 0,0\rangle$, intersected with $\overline{G}$. The resulting set consists of some part of the y-axis, and infinitely many segments of G, each cut off from each other. Such a set is not connected, and thus $\overline{G}$ does not contain arbitrarily small connected open neighbourhoods about $\langle 0,0\rangle$, as required.

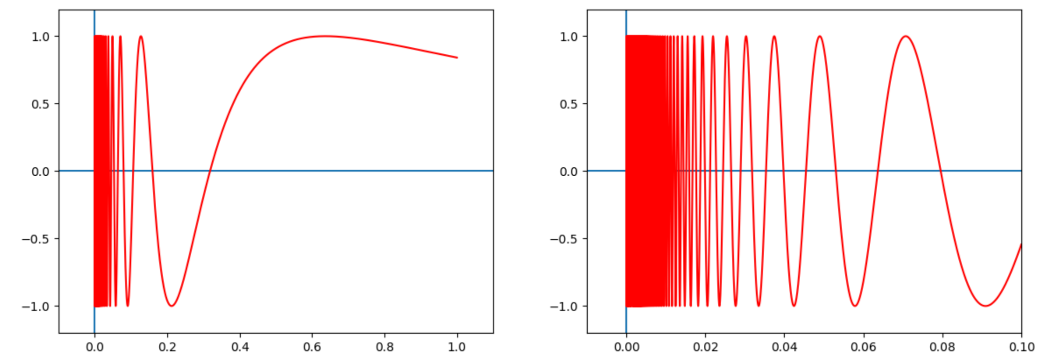

Figure 5: Left: The topologist's sine curve $\overline{G}$, shown in full on a section of $\mathbb{R}^2$. It consists of the graph of $\sin\frac{1}{x}$, along with the closed interval between -1 and 1 on the y-axis.

Right: We zoom in near the origin. The infinitely tight 'squeezing up' of the space towards the points that are not locally connected is typical of connected spaces that are not locally connected - see the 'infinite broom', or the cross-section of the cubic connectedness locus in Figure 6.

The local connectivity of the Mandelbrot set is not to be taken for granted. To demonstrate this, we give examples of compact, connected analogues to the Mandelbrot set for higher-degree polynomials, which are not locally connected.

Example 2.26 For each $a, b \in \mathbb{C}$, define the cubic polynomial $g_{a,b} : \mathbb{C} \to \mathbb{C}$ by $g_{a,b}(z) = z^3 - 3a^2z + b$. Similar to the quadratic case, these functions describe the dynamical behaviour of all complex cubics, since every cubic is affine conjugate to one of them. Also in parallel to the quadratic case, this is a good choice of functions to look at, since it has critical points a and $-a$, of a nice form (in particular, the critical points average to zero - a polynomial with this property is said to be *centred*). There is a subtlety though in that we don't quite have each cubic being affine conjugate to a unique $g_{a,b}$, since

[4]In fact, we have proven something more general. We used only the fact that $\overline{G}$ was the closure of a connected set, and so we may conclude that the closure of a connected set is always connected.

$g_{a,b}$ and $g_{a,-b}$ are affine conjugate. We quotient this out, defining $G = \{g_{a,b} : \langle a, b\rangle \in \mathbb{C} \times \mathbb{C}\}$, and considering the parameter space $G/\sim$, where $\sim$ identifies $g_{a,b}$ and $g_{a,-b}$. A typical element of the space is $[g] = \{g, g'\} \in G/\sim$, with g and g' having the same long term dynamical behaviour.

Recall the Dichotomy Theorem [2.17] for our family of complex quadratics. The result of Fatou and Julia was actually more general. We may define Julia sets in an analogous way for a cubic (or indeed any function), as the boundary of the set of initial points from which iterates stay bounded. Accordingly, the general result is that the Julia set is connected if (and only if) all of the critical points are bounded under iteration by the map in question, and that the Julia set is a Cantor set if all of the critical points go off to infinity under iteration. In the quadratic case, where there is only one critical point, a dichotomy is forced. But for our cubics, with two critical points, it may be the case that the two critical points have different dynamical behaviour - though even in this case, the Julia set is still not connected. Thus, our dichotomy is slightly weakened, with the Julia set connected if both critical points stay bounded, and disconnected otherwise. This suffices, though, to define the cubic connectedness locus

$$\begin{aligned}\mathcal{C}_3 &= \{[g] \in G/\sim\ :\ \text{The Julia set of } g \text{ is connected}\}\\ &= \{[g] \in G/\sim\ : |g^{\circ k}(a)| \not\to \infty \text{ and } |g^{\circ k}(-a)| \not\to \infty\}\end{aligned}$$

Branner and Hubbard [Branner and Hubbard, 1988] showed that $\mathcal{C}_3$ is compact and connected. Then, in his thesis, Pierre Lavaurs proved that the set was not locally connected [Lavaurs, 1989], confirming an earlier observation of Milnor from computer pictures. For a proof that is readily available online, see Appendix B of [Epstein and Yampolsky, 1999].

With a similar method, we may construct the connectedness locus $\mathcal{C}_d$ of the complex polynomials of a given degree $d \in \mathbb{N}$ for any $d \geq 2$, by starting with the space of all degree-d polynomials, and identifying all affine conjugate elements, and then taking the subset of conjugacy classes whose elements have a connected Julia set. For each natural number $d \geq 2$, it is possible to parametrise the affine conjugacy classes of the degree-d polynomials by $d-1$ complex variables, and in this way, we may view $\mathcal{C}_d$ as a subset of $\mathbb{C}^{d-1}$. Note that $\mathcal{C}_2$ is the Mandelbrot set. If $d \geq 3$, then $\mathcal{C}_d$ is not locally connected.[5]

Conjecture 1 (MLC) The Mandelbrot set M is locally connected.

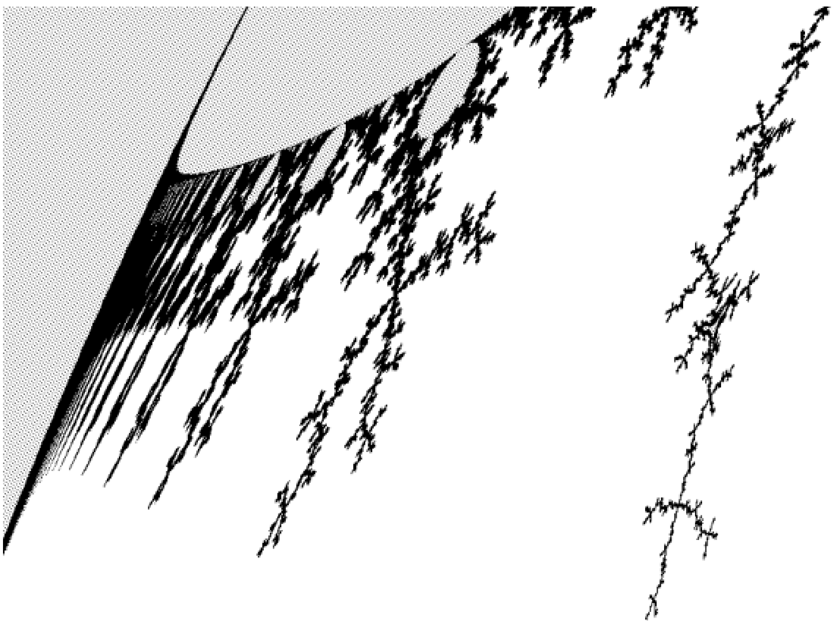

Figure 6: Some part of a cross-section of $\mathcal{C}_3$, where local connectivity fails in a similar way to Example 2.25. This image was taken from [Milnor, 1992], where it is Figure 19 on page 15.

[5]It is worth noting that the question of local connectivity of $\mathcal{C}_d$ is not as important for $d \neq 2$, because in these cases, it does not give us *density of hyperbolicity* in the corresponding parameter spaces (see Section 3.2 for exposition on density of hyperbolicity in the case $d = 2$).

3 An Introduction to Complex Dynamics Theory for Quadratic Maps

"To study complicated phenomena through their simplest incarnation; this is often the role of the mathematician."
—Adrien Douady

Before looking at the mathematics that has gone into attempts to obtain a proof of the conjecture, and the mathematics that would result from such a proof, we must build our theory a great deal. We introduce the fundamental ideas of the world of complex dynamics, specialised to our family F of quadratic maps. Note that a lot of the theory here can be generalised to other families of complex functions, most notably rational functions - for a more general treatment, see [Beardon, 2000], [Carleson and Gamelin, 1996] or [Milnor, 1990]. This section is based in large parts on the excellent article 'The Mandelbrot Set' by Bodil Branner in the book [Devaney and Keen, 1997].

3.1 Orbits, periodic points, and multipliers

The sequences of iterates $O(c,z) := \left(f_c^{\circ k}(z)\right)_{k\in\mathbb{N}}$ is called the *orbit* of $z \in \mathbb{C}$ under f_c. Note that for $k = 0$, we are taking $f_c^{\circ k}(z) = z$. Given Theorem 2.17 and its consequences for the family F, the orbit $O(c, 0)$ is of particular interest, and we will call it the *critical orbit.*

Points whose orbits repeat, that is $z \in \mathbb{C}$ such that $f_c^{\circ l}(z) = f_c^{\circ(l+m)}(z)$ for some $l \in \mathbb{N}, m \in \mathbb{Z}^+$ (and where we may freely take m to be the minimal such integer), are of particular interest, and are called *pre-periodic points.* If $l = 0$, the point is *periodic*, otherwise it is *strictly pre-periodic.* In this case the orbit is simply described, since

$$O(c,z) = \left(z, f(z), \ldots, f_c^{\circ l}(z), f_c^{\circ(l+1)}(z), \ldots f_c^{\circ(l+m-1)}(z), f_c^{\circ l}(z), f_c^{\circ(l+1)}(z), \ldots\right)$$

The sequence $\left(f_c^{\circ l}(z), f_c^{\circ(l+1)}(z), \ldots f_c^{\circ(l+m-1)}(z)\right)$ is called a *cycle*, and its length m is the *period* of the cycle. For a cycle of period 1, the solitary point of which it consists is called a *fixed point.* Note that every point in a cycle is necessarily periodic, and has the same orbit as the other points in its cycle (modulo a shifting of the sequence). Note also that every $f_c \in F$ has infinitely many cycles, since the Fundamental Theorem of Algebra guarantees roots to the polynomial equation $f_c^{\circ l}(z) = f_c^{\circ(l+m)}(z)$.

Consider the action of iterating points in $\mathbb{C}$ by the m-fold composition of f_c. We obtain a function $f_c^{\circ m} : \mathbb{C} \to \mathbb{C}$. Let $(z_1, \ldots, z_m)$ be a period m cycle of f_c. The Taylor series of $f_c^{\circ m}$ at a periodic point gives us information about the iterative behaviour of nearby points. In particular, let $r > 0$ be sufficiently small, and let $t \in B(z_1, r)$ be a point in the plane 'near' the periodic point z_1. Then

$$\begin{aligned} f_c^{\circ m}(t) &= f_c^{\circ m}(z_1) + \left(f_c^{\circ m}\right)'(z_1) \cdot (t - z_1) + [(t - z_1)^2 \text{ and higher order terms}] \\ &\approx z_1 + \left(f_c^{\circ m}\right)'(z_1) \cdot (t - z_1) \end{aligned}$$

Thus, if $\left|\left(f_c^{\circ m}\right)'(z_1)\right|$ is smaller than 1, $f_c^{\circ m}(t)$ is closer to z_1 than t was, and $f_c^{\circ 2m}$ closer still, etc. If instead $\left|\left(f_c^{\circ m}\right)'(z_1)\right|$ is larger than 1, t is 'pushed away' from z_1, with each iterate of $f_c^{\circ m}$ going further and further away. We feel sufficiently motivated to make a definition.

Definition 3.1 The *multiplier* of a periodic point z of period m is $\rho := \left(f_c^{\circ m}\right)'(z)$.

In fact, we will speak of the multiplier of a cycle, as opposed to of a particular periodic point, since the multipliers of distinct points in a cycle are equal. This is due to the following result, which allows for easy calculation of the multiplier.

Proposition 3.2 *Let $(z_1, \ldots, z_m)$ be a cycle of $f_c \in F$, and let z_1 have multiplier ρ. Then*

$$\rho = \prod_{k=1}^{m} f_c'(z_k) = 2^m \prod_{k=1}^{m} z_k$$

Proof. By the chain rule,

$$\begin{aligned}(f_c^{\circ m})' &= \left(f_c \circ f_c^{\circ(m-1)}\right)' = \left(f_c' \circ f_c^{\circ(m-1)}\right) \cdot \left(f_c^{\circ(m-1)}\right)' \\ &= \ldots = \left(f_c' \circ f_c^{\circ(m-1)}\right) \cdot \left(f_c' \circ f_c^{\circ(m-2)}\right) \cdot \ldots \cdot \left(f_c' \circ f_c\right) \cdot f_c'\end{aligned}$$

Therefore,

$$\begin{aligned}(f_c^{\circ m})'(z_1) &= \left(f_c'\left(f_c^{\circ(m-1)}(z_1)\right)\right) \cdot \ldots \cdot f_c'(z_1) \\ &= f_c'(z_m) \cdot \ldots \cdot f_c'(z_1)\end{aligned}$$

Up to this point, the proof is applicable to arbitrary functions. Since we are working in particular with the family of maps F, we may go further. For any $c, z \in \mathbb{C}$, $f_c'(z) = 2z$, so that

$$(f_c^{\circ m})'(z_1) = \prod_{k=1}^{m}(2z_k) = 2^m \prod_{k=1}^{m} z_k$$

□

Recalling our earlier observations that nearby points are sucked in or pushed out depending on the size of the multiplier, and since all points in a cycle have the same multiplier, we are able to classify cycles accordingly.

Definition 3.3 Let $(z_1, \ldots z_m)$ be a cycle of $f_c \in F$ that has multiplier ρ. Then the cycle, and the periodic points $z_1, ..., z_m$, are

(i) *attracting* if $|\rho| < 1$,

(ii) *indifferent* if $|\rho| = 1$, or

(iii) *repelling* if $|\rho| > 1$.

Further, if $\rho = 0$, we say the cycle and its points are *super-attracting*. By Proposition 3.2, a cycle $(z_1, ..., z_m)$ is super-attracting if and only if $z_j = 0$ for some $1 \leq j \leq m$.

Indifferent cycles have $|\rho| = 1$, so that $\rho = e^{2\pi i t}$ for some $t \in \mathbb{R}/\mathbb{Z}$ (i.e. the set of points on the unit circle, parametrized from 0 to 1). If $t \in \mathbb{Q}/\mathbb{Z}$, we say the cycle is *rationally indifferent* or *parabolic*, and if $t \notin \mathbb{Q}/\mathbb{Z}$, we say the cycle is *irrationally indifferent*. The special case of a parabolic cycle with multiplier exactly 1 is called a *primitive* parabolic cycle.

Example 3.4 Recall Example 2.2. We considered a solution $\omega \approx -0.12+0.745i$ to $(c^2+c)^2+c = 0$, so that $O(\omega, 0) = \left(0, \omega, \omega^2 + \omega, 0, ...\right)$. Thus, we have a cycle $(0, \omega, \omega^2 + \omega)$, which in particular contains the critical point $z = 0$, and so must be super-attracting. In this context, Figure 1 makes sense. Starting at $z = 0.2$, the iterates are attracted towards and converge to our length 3 cycle. Note that for the below picture, we have actually perturbed the parameter as well, since ω is irrational. We will soon discuss hyperbolic components, and how slight perturbations of parameters with attracting cycles gives only slight perturbations of orbits.

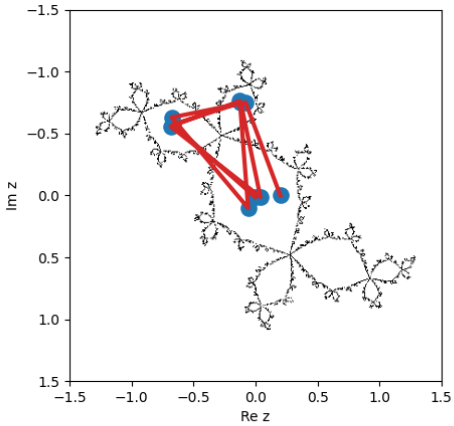

Figure 7: We plot $O(\omega, 0.2)$, overlayed on the Julia set $J(\omega)$. The orbit is pulled towards the super-attracting cycle $\left(0, \omega, \omega^2 + \omega\right)$.

Example 3.5 Consider the point $c = \frac{1}{4}$. From the proof of Theorem 2.7, $c \in M$. Also, $f_c\left(\frac{1}{2}\right) = \left(\frac{1}{2}\right)^2 + \frac{1}{4} = \frac{1}{4} + \frac{1}{4} = \frac{1}{2}$, so that

$$O\left(c, \frac{1}{2}\right) = \left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}, \dots\right)$$

and we have a fixed point of f_c. Its multiplier is $f_c'\left(\frac{1}{2}\right) = 2\left(\frac{1}{2}\right) = 1$, so that in fact $\frac{1}{2}$ is a primitive parabolic fixed point. The parameter c is a 'messy' point of M, in part because of the multiplier of its fixed point. Calculating and generating images of its Julia set is more difficult than usual, since it contains the indifferent point, around which iterates move *very slowly*. One popular method of generating images of Julia sets is to 'iterate backwards' by taking pre-images of some repelling periodic point on the Julia set. This necessarily fills out the Julia set (in that the resulting set of points is dense in the Julia set), but will be particularly slow if we were to apply it to $J(\frac{1}{4})$, since progress will be slowed whenever we are near to $z = \frac{1}{2}$.

We note also that c is on the boundary of M. Certainly there exist points of M in any neighbourhood of c, since $B(0, \frac{1}{4}) \subseteq M$. We claim that for real $c' > \frac{1}{4}$, we necessarily have $c' \notin M$, since for any $k \in \mathbb{Z}^+$

$$f_{c'}^{\circ k}(0) - f_{c'}^{\circ(k-1)}(0) = f_{c'}\left(f_{c'}^{\circ(k-1)}(0)\right) - f_{c'}^{\circ(k-1)}(0) = \left(f_{c'}^{\circ(k-1)}(0)\right)^2 + c' - f_{c'}^{\circ(k-1)}(0) \geq c' - \frac{1}{4}$$

and therefore

$$\begin{aligned} f_{c'}^{\circ k}(0) &\geq f_{c'}^{\circ(k-1)}(0) + \left(c' - \frac{1}{4}\right) \\ &\geq f_{c'}^{\circ(k-2)}(0) + 2\left(c' - \frac{1}{4}\right) \\ &\vdots \\ &\geq f_{c'}(0) + (k-1)\left(c' - \frac{1}{4}\right) \end{aligned}$$

Thus, given that $c' - \frac{1}{4} > 0$, we have $f_{c'}^{\circ k}(0) \to \infty$ as $k \to \infty$, and we may conclude that $\left|f_{c'}^{\circ k}(0)\right| \to \infty$ as $k \to \infty$, as required.

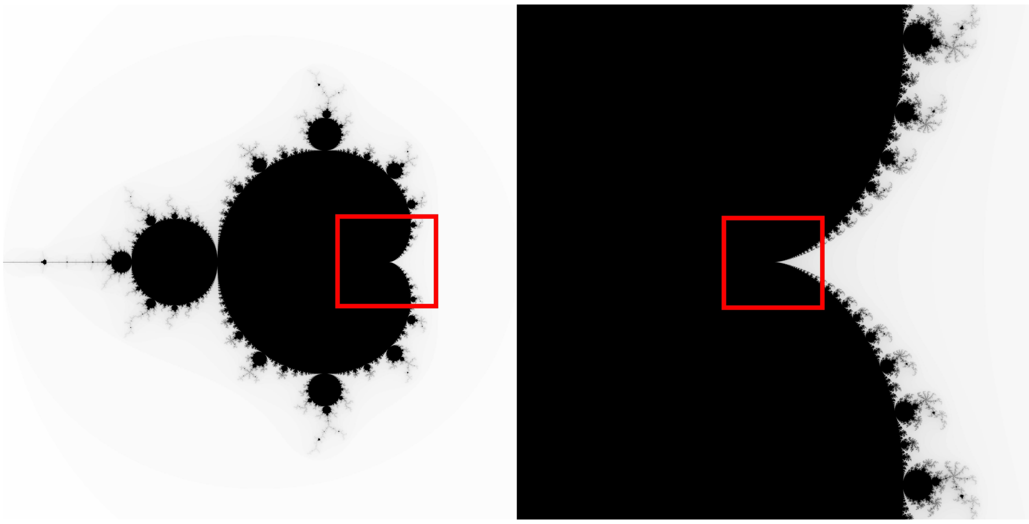

Figure 8: The 'cusp' $c = \frac{1}{4}$ of the Mandelbrot set. It is the root of the main cardioid (See Section 3.2). Given the infinity of limbs collapsing upon it, it can be a harsh parameter to get a handle on. Thankfully, we do have local connectivity here - see Theorem 4.24.

As a quick aside, we are fairly close to a full description of $M \cap \mathbb{R}$. We have that $|c| > 2 \Rightarrow c \notin M$ by Corollary 2.6, that $B(0, \frac{1}{4}) \subseteq M$ by Theorem 2.7, and now that $c > \frac{1}{4} \Rightarrow c \notin M$. Since $O(-2, 0) = (0, -2, 2, 2, ...)$ is clearly a bounded orbit and so $-2 \in M$, we might guess that $M \cap \mathbb{R} = [-2, \frac{1}{4}]$, and indeed this may be shown to be the case.

We end the section by looking at a class of parameters that are easily defined, given the significant role they play in the study of M.

Definition 3.6 Suppose $f_c \in F$ is such that the critical point $z = 0$ is strictly pre-periodic, that is, there is some $l, m \in \mathbb{Z}^+$ such that $f_c^{\circ l}(0) = f_c^{\circ(l+m)}(0)$, and $f_c^{\circ k}(0) = 0$ if and only if $k = 0$. Then c is called a *Misiurewicz point*.

Example 3.7 We observed earlier that $M \cap \mathbb{R} = [-2, \frac{1}{4}]$, and also that $O(-2, 0) = (0, -2, 2, 2, \ldots)$. Therefore, -2 lies on the boundary of M, and is a Misiurewicz point. It may be shown that the filled-in Julia set is $K(-2) = [-2, 2]$. This has no interior, and is its own boundary in $\mathbb{C}$, so that $J(-2) = K(-2)$.

Example 3.8 Another simple example of a Misiurewicz point is $c = i$, since $O(i, 0) = (0, i, -1 + i, -i, -1 + i, ...)$. What does the Julia set look like?

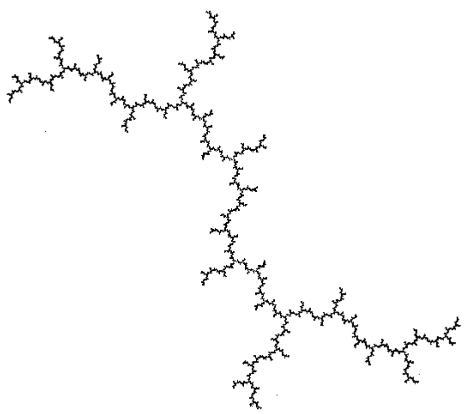

Figure 9: The filled-in Julia set $K(i)$, which has no interior and is equal to its boundary, so that $J(i) = K(i)$.

The previous examples are no coincidence. It may be shown that, for a Misiurewicz parameter c, we must have $J(c) = K(c)$. In this case, the Julia set is called a *dendrite.*

It is clear from the definition that all Misiurewicz points belong to M. In fact, Douady and Hubbard showed that Misiurewicz points must lie on the boundary of M. They showed also that these points are dense on the boundary, and that the cycle on to which 0 exactly falls is necessarily repelling (see [Douady and Hubbard, 1985a]).

The following is certainly worth noting. It was first proved by Yoccoz, and is first stated explicitly in the literature in [Lei, 1992]. For an english-language proof, see Theorem 6.4 of [Schleicher, 1999].

Theorem 3.9 *Let $c \in M$ be a Misiurewicz parameter. Then M is locally connected at c.*

We make one last remark regarding Misiurewicz points. The geometric similarities between Julia sets and the Mandelbrot set are remarkable and well-known. At Misiurewicz points, these similarities are at their most pronounced. Let $c \in M$ be a Misiurewicz point. Tan Lei [Lei, 1990] proved that M and $J(c)$ are both self-similar as we zoom into $c \in \mathbb{C}$ in either set. Furthermore, the neighbourhoods of c in M and $J(c)$ can (in a precise sense) be made to look arbitrarily similar to each other, if we zoom in enough. Note that since $J(c) = K(c)$, we have by definition that $0 \in J(c)$, so that $c \in J(c)$ by Theorem 2.10, as you would hope given the statement of this paragraph!

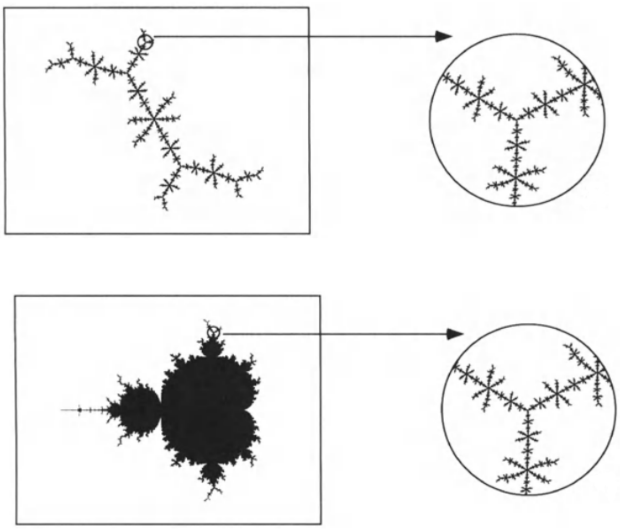

Figure 10: The point $c \approx -0.1011 + 0.9563i$ is a Misiurewicz point. At the top, we have an image of the dendrite $J(c)$, with a zoomed picture of a neighbourhood of c. On the bottom, we have M, and a zoomed picture of a neighbourhood of c. This picture is taken from page 887 of the book [Peitgen et al., 1992].

3.2 Hyperbolicity and the Douady-Hubbard-Sullivan Theorem

Recall that for any $f_c \in F$, there are infinitely many cycles. However, the family F in particular is very restricted in the number of attracting cycles its elements can possess. This is a consequence of the following theorem, proved by Fatou in 1905, which will lead us to a discussion of hyperbolicity and its relationship with the MLC conjecture. Note also that this theorem reinforces the idea that the critical orbit has a particular significance.

Theorem 3.10 *Let $f_c \in F$, and let $(\alpha_1, \ldots, \alpha_m)$ be any attracting cycle of f_c. Then the critical point $z = 0$ is attracted towards the cycle, i.e. there exists $l \in \mathbb{N}$ such that*

$$\lim_{k\to\infty} \left(f_c^{\circ(l+n+km)}(0) \right) = \alpha_n$$

for each $n = 1, ..., m$.

We record a couple of the immediate consequences of Theorem 3.10.

Corollary 3.11 *Let $f_c \in F$. Then f_c has at most one attracting cycle.*

Proof. By Theorem 3.10, the critical point is attracted towards all attracting cycles of f_c. If there were two or more attracting cycles, the critical point would have to tend to both, which is clearly impossible. □

Corollary 3.12 *Let $f_c \in F$ possess an attracting cycle, say $(\alpha_1, \ldots, \alpha_m)$. Then $c \in M$.*

Proof. By Theorem 3.10, there is some $k \in \mathbb{N}$ such that $\left|f_c^{\circ l}(0)\right| < \max\{|\alpha_1| + 1, ..., |\alpha_m| + 1\}$ for all $l > k$, and from this bound we may conclude $c \in M$. □

In the more general context of other families of functions, we say a map is *hyperbolic* if all critical points tend to an attracting cycle under iteration. This notion is equivalent to other definitions of hyperbolicity appearing in the literature, if we extend the dynamical space to the Riemann sphere $\widehat{\mathbb{C}} := \mathbb{C} \cup \{\infty\}$, and define $f_c(\infty) = \infty$ to be an attracting fixed point[6]. Then, a parameter c and its associated quadratic are hyperbolic if and only if $c \notin M$, or $c \in M$ and f_c has an attracting cycle.

For over a century, dynamicists have been trying to answer the following question, which may also be asked for other classes of maps such as the 'Multibrot' families $\{f : z \mapsto z^d + c \mid c \in \mathbb{C}\}$ for a fixed integer $d > 2$ and the family of exponentials $\{f : z \mapsto e^z + c \mid c \in \mathbb{C}\}$, but which we ask only for the family F:

Conjecture 2 (Density of Hyperbolicity for complex quadratic maps) Hyperbolic parameters are dense in the paramater plane of the family F. That is, for each $c \in \mathbb{C}$, and for all $\varepsilon > 0$, there exists a parameter $c' \in B(c, \varepsilon)$ such that $f_{c'}$ is a hyperbolic map.

Generally, the dynamics of hyperbolic maps are well-understood. Thus, if the hyperbolic maps are dense in a given space, we can approximate every map arbitrarily well by a hyperbolic map, whose dynamics we understand. Chapter 2 of [McMullen, 1994b] has a good, accessible exposition of this. Chapter 14 of [Milnor, 1990] is, meanwhile, more in-depth.

Since all parameters in the complement of M are hyperbolic, we focus in specifically on elements of M that are hyperbolic.

By Theorem 3.10 and its corollaries, we may classify these parameters by the period of their unique attracting cycle. The most simple case, then, is that of parameters whose quadratics possess an attracting fixed point. Define $W_1 = \{c \in \mathbb{C} : f_c \text{ has an attracting fixed point}\}$. It is easy to proceed

[6] At infinity, we define the multiplier as the reciprocal of the usual definition. We should also point out that ∞ is excluded from consideration in Corollary 3.11.

algebraically. Let z be the attracting fixed point. We must have $z^2 + c = z$, with the multiplier $|\rho| = |2z| < 1$. Then we must have

$$z = \frac{1 \pm \sqrt{1-4c}}{2} \Rightarrow |\rho| = |2z| = \left|1 \pm \sqrt{1-4c}\right| < 1$$

and we must necessarily take the negative root, since $\sqrt{1-4c} > 0$ (taking the principal root). Conversely, any such c satisfying the inequality $\left|1 - \sqrt{1-4c}\right| < 1$ must be a parameter with an attracting fixed point, since we may then use our formula to find the fixed point z, and by definition its multiplier $\left|1 - \sqrt{1-4c}\right|$ makes it attracting. Therefore, $W_1 = \{c \in \mathbb{C} : \left|1 - \sqrt{1-4c}\right| < 1\}$. Refer to Figure 11 for a look at W_1, and its uncanny familiarity... Note the boundary of W_1 consists precisely of the parameters that possess indifferent fixed points, an example of which being $c = \frac{1}{4}$ as in Example 3.5. Note also that W_1 does therefore not contain any of its boundary points, so is open.

The next most simple case is those parameters whose quadratics possess an attracting cycle of period 2. Define $W_2 = \{c \in \mathbb{C} : f_c \text{ has an attracting period 2 cycle}\}$. Again, we are able to proceed algebraically. Let $(z, f_c(z))$ be the attracting cycle. We require

$$\begin{aligned} f_c^{\circ 2}(z) = z &\iff f_c(z^2 + c) = (z^2 + c)^2 + c = z \\ &\iff z^4 + 2cz^2 - z + c^2 + c = 0 \\ &\iff (z^2 - z + c)(z^2 + z + c + 1) = 0 \end{aligned}$$

But $z^2 - z + c = 0$ if and only if $z^2 + c = z$, if and only if z is a fixed point. Thus, c has an attracting cycle $(z, f_c(z))$ if and only if $z^2 + z + c + 1 = 0$, with the multiplier ρ satisfying $|\rho| < 1$. We proceed using Proposition 3.2 to calculate the multiplier, and then perform suitable algebraic rearranging and substituting using our quadratic equation:

$$\rho = 4z f_c(z) = 4z(z^2 + c) = 4z(-z - 1) = -4(z^2 + z) = 4(c + 1)$$

Thus, the inequality that describes W_2 is $|4(c + 1)| < 1$. Letting $c = x + yi$ and rearranging, we obtain

$$(x + 1)^2 + y^2 < \left(\frac{1}{4}\right)^2$$

so that $W_2 = B\left(-1, \frac{1}{4}\right)$. Again, note that the boundary consists of those points with $\rho = 1$, and that our set is open.

Consider these two subsets of the interior of M that we have located. They are disjoint, but they 'touch' at a single point, in that $\overline{W_1} \cap \overline{W_2} = \{-\frac{3}{4}\}$. As a result, $c = -\frac{3}{4}$ may be described as a *bifurcation point* of M. If we were to start in W_1, and move c over in a continuous fashion towards W_2, we would find the multiplier of the attracting fixed point gradually tend towards 1 (in this sense, the point is becoming 'less' attracting), and at the point of crossing $c = -\frac{3}{4}$, this fixed point becomes indifferent. As we enter W_2, the fixed point will now be a repelling one, while an attracting cycle of period 2 has now appeared.[7]

All of these observations hold in a far more general sense. In fact, for any $k \in \mathbb{Z}^+$, we may find open sets (in general, we get multiple disjoint open sets for a given k) consisting precisely of all parameters whose quadratics possess an attracting cycle of period k. Any part of M with a non-empty interior that can be seen in Figure 3 with the naked eye is one of these open sets. Further, they are all disjoint from one another, so that these subsets we describe are (at least some of) the connected components of $\text{int}(M)$, and they are accordingly named the *hyperbolic components* of M. Note that we would not be able to 'walk' between all hyperbolic components freely, via bifurcation points, since some hyperbolic components are within mini-Mandelbrot copies, as in Figure 3.

[7]This example is the first step in the Mandelbrot set's version of a universal phenomenon of bifurcation theory, closely related to the famous *Feigenbaum constant*. We return to this in Example 4.16.

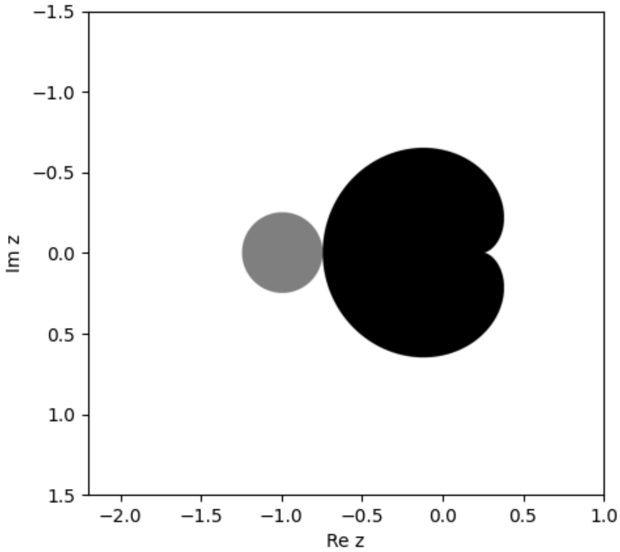

Figure 11: The hyperbolic component W_1, consisting of precisely those parameters whose quadratic possesses an attracting fixed point, forms the main cardioid of the Mandelbrot set, and is shown in black. Similarly, the hyperbolic component W_2 for period 2 cycles is shown in grey.

We again draw attention to the fact that the hyperbolic components W_1 and W_2 are special cases, in that there is a single hyperbolic component for the parameters with attracting cycles with period 1 or 2. This is the case only in these two instances. For example, there are 3 hyperbolic components for period 3 attracting cycles, and 65535 hyperbolic components for period 17 attracting cycles. This sequence is recorded in the Online Encyclopedia of Integer Sequences as A000740.

We have in fact made progress in our great endeavour of showing local connectivity at all points of M!

Theorem 3.13 *Let $c \in M$ be hyperbolic. Then M is locally connected at c.*

Proof. Our supposition is equivalent to f_c having some attracting cycle, say $(\alpha_1, \ldots, \alpha_m)$. An application of the *Implicit Function Theorem* to the function $G(z, c) = f_c^{\circ m}(z) - z$ gives that there is some $r > 0$ such that every parameter in $B(c, r)$ has an attracting cycle. Thus, $B(c, r) \subseteq M$, and $c \in \text{int}\,(M)$. It follows that M is locally connected at c, essentially because $\mathbb{C}$ itself is locally connected. In particular, given any open set U in M that contains c, we can find an open ball in $\mathbb{C}$ centred at c which is entirely contained in U and thus M, and so is certainly a connected neighbourhood of c in M. □

We will later show (Section 4.3) that M is locally connected at all points on the boundary of a hyperbolic component. In the specific cases of W_1 and W_2, we saw that such points necessarily had indifferent cycles. Our next step is to see that this holds for all hyperbolic components.

By the groundbreaking work of A. Douady, J. Hubbard and D. Sullivan in the 1980s, we may use the multiplier of the attracting cycles in a hyperbolic component as a sort of roadmap of the component. First, a definition is needed - conformal isomorphisms are an important type of mapping that preserve some structure.

Definition 3.14 Let $D, E \subseteq \mathbb{C}$ be non-empty, connected and open. A *conformal mapping* $f : D \to E$ is a mapping which 'preserves angles', in that any two intersecting curves drawn in D will have their angle of intersection preserved by f. A *conformal isomorphism* is a conformal mapping which is also a bijection. If there is a conformal isomorphism $D \to E$, we say D and E are *conformally isomorphic* or *conformally equivalent*.

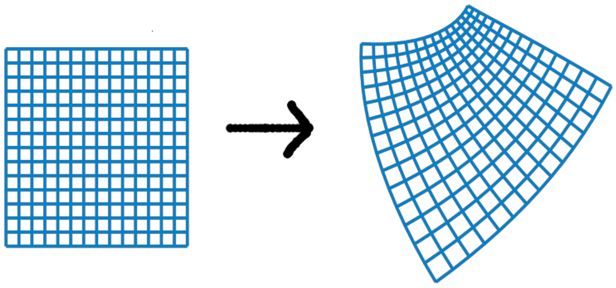

Figure 12: An illustration of a conformal mapping. Although the space is distorted, angles between curves and lines are preserved by the mapping. This is demonstrated by the perpendicular gridlines, which are still at right angles under the mapping. Picture taken from the Wikipedia article on conformal mappings, and released into the public domain by its creator Oleg Alexandrov.

It may be shown that conformal equivalence is an equivalence relation. In Section 4.3, we will develop a small bit of theory behind conformal mappings and isomorphisms - this will contextualise the below result as well as the occasional reference we will make to these mappings throughout this section.

Theorem 3.15 (Douady-Hubbard-Sullivan Theorem) *Let W be a hyperbolic component of M, and define the map $\rho_W : W \to B(0,1)$ by $\rho_W(c) = \rho_c$, where ρ_c is the multiplier of the unique attracting cycle of $c \in W$. Then ρ_W is a conformal isomorphism. The map can be extended to a homeomorphism $\overline{W} \to D(0,1)$.*

At first, it seems that there is a difficulty in extending the map to the boundary, since the points on the boundary do not have a unique attracting cycle (or any attracting cycle at all) whose multiplier we can use. This is circumnavigated in the following way: As we approach a boundary point $c \in \mathbb{C}$ via nearby points in W, the unique attracting cycle of period k 'varies analytically', in that small perturbations of the parameter result in small perturbations of the locations of the k points in the cycle, as well as a small perturbation of the multiplier of the cycle. This behaviour persists even as we hit the boundary, so that we can take the limit of the multiplier as we approach c. As a result, since ρ_W is approaching the boundary of the unit ball, we will have $\rho_W(c) = e^{2\pi it}$ for some $t \in \mathbb{R}/\mathbb{Z}$, and f_c will have an indifferent cycle of period k. The only exception to this is when $t = 0$ (equivalently $\rho = 1$). This point of the hyperbolic component is called the *root*. There are two possibilities for a root c of a hyperbolic component H, of period k: 1) It is the bifurcation point between two hyperbolic components. In this case, H touches another hyperbolic component G, whose period m must divide k. We will have $\overline{H} \cap \overline{G} = \{c\}$. As we approach c in H, points in the cycle partition themselves, with each part of the partition consisting of points tending to a common limit. At c, each of these parts coalesce into a single point so that the attracting cycle of k points becomes an indifferent cycle of m points. 2) It is the cusp $c = 1/4$ of the main cardioid (see Example 3.5), or is the cusp of the cardioid of some miniature homeomorphic copy of M within itself (see Section 4.2).

As a result, we are provided with a parametrization of the boundary $\gamma_W : \mathbb{R}/\mathbb{Z} \to \partial W$ given by $\gamma_W(t) = \rho_W^{-1}(e^{2\pi it})$. For a general boundary point $\gamma_W(t) = \rho_W^{-1}(e^{2\pi it})$ of a given hyperbolic component W, we will say $\gamma_W(t)$ has *internal angle* t. The *centre* of W is the unique parameter in W whose quadratic possesses a super-attracting cycle, that is $\rho_W^{-1}(0)$. Given this is the point that the Douady-Hubbard-Sullivan map identifies with the centre of the unit disc, our nomenclature makes sense.

A further consequence of Theorem 3.15 is that each hyperbolic component contains precisely one parameter whose attracting cycle has a given multiplier ρ (for $|\rho| < 1$). Therefore, the number of parameters in M whose quadratics possess an attracting cycle of a given period with a given multiplier is exactly the number of hyperbolic components that consist of parameters with attracting cycles of that period, so is given by the sequence A000740.

It seems as though the hyperbolic components are filling the interior of M, so it is natural to next ask for a full description of the interior of M, that is, what are the maps with bounded critical orbit for which all nearby maps also have bounded critical orbit? There is a deep link, in fact an equivalence, between this question and Conjecture 2. For, we have seen that all hyperbolic parameters that are in M are in its interior. If the interior could be shown to consist only of hyperbolic parameters, this would yield density of hyperbolicity, due to the following result.

Proposition 3.16 *The interior of M is dense in M, i.e.* $\overline{\text{int}(M)} = M$.

Proof. See Section VIII, Theorem 1.5 of [Carleson and Gamelin, 1996]. □

Suppose instead that we find a non-hyperbolic parameter $c \in \text{int}(M)$. Since it is in the interior of M, there is some $\varepsilon > 0$ such that $B(c, \varepsilon) \subseteq \text{int}(M)$. This ball cannot contain any hyperbolic parameters, else it would be contained in a hyperbolic component, contradicting non-hyperbolicity of c. Then our parameter c does not have arbitrarily close hyperbolic parameters, and we do not have density of hyperbolicity in the space of complex quadratics. We conclude:

Theorem 3.17 *Conjecture 2 is equivalent to the statement*

$$\text{int}(M) = \{c \in M : f_c \text{ possesses an attracting cycle}\}$$

As outlined above, any non-hyperbolic parameters of $\text{int}(M)$ would have to live inside a connected component of $\text{int}(M)$ consisting entirely of non-hyperbolic parameters. No such component has ever been found. Given their elusive, hypothetical nature, these components are called *ghost components*. By Theorem 3.18 below, finding a ghost component would disprove the MLC conjecture. It may be shown that all parameters whose quadratics possess indifferent cycles belong to the boundary of M, so that any ghost component would consist of parameters whose cycles were all repelling.

The stated goal, and crowning achievement, of the acclaimed work [Douady and Hubbard, 1985a] is to assume local connectivity of M and use this to disprove the existence of ghost components of M. In this way, DHC was precisely the original motivation for the MLC conjecture.

Theorem 3.18 *The MLC conjecture, if true, would imply density of hyperbolicity for complex quadratic maps.*

A proof of Theorem 3.18 will be sketched in Section 5.4.

3.3 The Böttcher mapping, equipotentials and external rays

Here, we will use the power of complex analysis and some of its biggest theorems to introduce the most fundamental objects in tackling the MLC conjecture. We build up to a new characterisation of local connectivity for M, and in doing so construct a parametrization of its boundary (assuming MLC).

The definitions we will eventually make in this section are in a sense motivated by the below result.

Proposition 3.19 *Let $K \subseteq \mathbb{C}$ be compact. Then there exists a unique function $G_K : \mathbb{C} \backslash K \to \mathbb{R}$ satisfying the following properties:*

(i) G_K is harmonic, that is we may write $G_K(x + iy) = u(x, y) + iv(x, y)$, where $u, v : \mathbb{R}^2 \to \mathbb{R}$ satisfy $\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} = 0$,

(ii) $G_K(z) \to \infty$ as $z \to \infty$, and more specifically, $G_K(z) \sim \log |z|$ as $z \to \infty$, and

(iii) $G_K(z) \to 0$ as $z \to K$ (by which we mean, as $\min\{|z - t| : t \in K\} \to 0$).

The function G_K is called the *Green's function* of K.

Now, recall that the filled-in Julia set $K(c)$ of any $f_c \in F$ is compact. Then, we may define the function $G_c : \mathbb{C}\backslash K(c) \to \mathbb{R}$ to be the unique Green's function of $K(c)$. We may continuously extend the domain to all of $\mathbb{C}$ by setting $G_c(z) = 0$ for all $z \in K(c)$.

It is, perhaps surprisingly, not difficult to obtain an explicit formula for (the extension of) G_c.

Proposition 3.20 *Define* $\log^+ : [0, \infty) \to \mathbb{R}$ *by*

$$\log^+(x) = \begin{cases} 0 & \text{if } 0 \leq x \leq 1 \\ \log(x) & \text{otherwise} \end{cases}$$

Then

$$G_c(z) = \lim_{k \to \infty} \frac{\log^+ \left| f_c^{\circ k}(z) \right|}{2^k}$$

The function G_c defined above tells us how quickly iterates of f_c are escaping to infinity at each point $z \in \mathbb{C}$.

There is a related function $\phi_c : \mathbb{C}\backslash K(c) \to \mathbb{C}\backslash D(0,1)$ which is more useful to us. Its definition begins by an application of the following theorem.

Theorem 3.21 (Böttcher's Theorem) [Milnor, 1990]
Let $f(z) = a_n z^n + a_{n+1} z^{n+1} + \ldots$, where $n \geq 2, a_n \neq 0$. Then there is a local holomorphic change of coordinates $w = \phi(z)$ which conjugates f to the map $w \mapsto w^n$ throughout some neighbourhood of $\phi(0) = 0$.

To use the theorem, we must extend f_c to the Riemann sphere $\mathbb{C} \cup \{\infty\}$, and consider f_c in a neighbourhood of ∞. Equivalently, consider $f_c(1/z)$ in a neighbourhood of 0. The Laurent series for this function, appropriately translated, brings us to the setup of Theorem 3.21 with $n = 2$. Therefore, for any $c \in \mathbb{C}$, there exists a function ϕ_c defined in a neighbourhood of infinity, which conjugates f_c to f_0. Douady and Hubbard [Douady and Hubbard, 1985a] showed that if $c \in M$, then ϕ_c can be extended in a unique way to a conformal isomorphism $\phi_c : \mathbb{C}\backslash K(c) \to \mathbb{C}\backslash D(0,1)$, and in a way that maintains our conjugation. In fact, ϕ_c is the unique conformal isomorphism $\mathbb{C}\backslash K(c) \to \mathbb{C}\backslash D(0,1)$ that conjugates f_c to f_0. Thus, we have

$$\phi_c^{-1} \circ f_c \circ \phi_c = f_0 \Rightarrow \phi_c\left(f_c(z)\right) = f_0\left(\phi_c(z)\right) = \phi_c(z)^2$$

for all $z \in \mathbb{C}\backslash K(c)$.

Even if $c \notin M$, we can still get somewhat of a handle on things. In particular, it may be shown that c itself is in the domain of ϕ_c for all $c \in \mathbb{C}$. This will be crucial shortly, when we look instead at the parameter plane.

We call ϕ_c the *Böttcher mapping* of f_c, and it is related to the Green's function G_c in the following way.

Proposition 3.22 *Let $c \in \mathbb{C}$, and let $z \in \mathbb{C}\backslash K(c)$. Then*

$$G_c(z) = \log |\phi_c(z)|$$

Corollary 3.23 *$\phi_c(z) \sim z$ as $|z| \to \infty$.*

Proof. Follows immediately from Proposition 3.22 and *(ii)* of Proposition 3.19. □

We are now in a position to define probably the two most important objects regarding the MLC conjecture. Consider ϕ_c^{-1}. Its domain is $\mathbb{C}\backslash D(0,1) = \{z \in \mathbb{C} : |z| > 1\} = \{re^{2\pi it} : r > 1, t \in \mathbb{R}/\mathbb{Z}\}$.

Definition 3.24 Fix $r > 1$. The *dynamical equipotential* of $f_c \in F$, of potential r, is the set

$$\Gamma_c(r) = \{\phi_c^{-1}\left(re^{2\pi it}\right) : t \in \mathbb{R}/\mathbb{Z}\}$$

It is a closed, simple curve that loops around $K(c)$. As $r \to 1$, this loop is pulled tighter and tighter around the boundary of $K(c)$. The equipotential is the curve of all points of the dynamical plane at which the iterates of f_c are escaping to infinity at the same rate. Indeed, from the definition, we see that $\Gamma_c(r) = G_c^{-1}(\log r)$, so that the equipotential consists of points that all have the same fixed value under the Green's function.

Proposition 3.25 *Let $f_c \in F$. Then $f_c(\Gamma_c(r)) = \Gamma_c(r^2)$.*

Proof. Recall that $\phi_c(f_c(z)) = \phi_c(z)^2$ for all $z \in \mathbb{C}\backslash K(c)$. In particular, this relation will hold for all points belonging to some equipotential, since all equipotentials are disjoint from $K(c)$. Next, we have the sequence of equivalences

$$\begin{aligned}
z \in \Gamma_c(r) &\iff z = \phi_c^{-1}(re^{2\pi it}) \text{ for some } t \in \mathbb{R}/\mathbb{Z} \\
&\iff \phi_c(z) = re^{2\pi it} \text{ for some } t \in \mathbb{R}/\mathbb{Z} \\
&\iff \phi_c(z)^2 = r^2e^{2\pi i(2t)} = \phi_c(f_c(z)) \text{ for some } t \in \mathbb{R}/\mathbb{Z} \\
&\iff \phi_c(f_c(z)) = r^2e^{2\pi it} \text{ for some } t \in \mathbb{R}/\mathbb{Z} \\
&\iff f_c(z) = \phi_c^{-1}(r^2e^{2\pi it}) \text{ for some } t \in \mathbb{R}/\mathbb{Z} \\
&\iff f_c(z) \in \Gamma_c(r^2)
\end{aligned}$$

where we have used the aforementioned relation, as well as the fact that ϕ_c is bijective. Note also that the t in question need not stay fixed as we go from any one statement to the next or last one. For example, if $\phi_c(z)^2 = r^2e^{2\pi i(2t)}$ for some $t \in \mathbb{R}/\mathbb{Z}$, then we may in fact have $\phi_c(z) = -re^{2\pi it} = re^{\pi i}e^{2\pi it} = re^{2\pi i(t+1/2)}$. But then taking $t+1/2$ as our 'new t', the implication stated above still holds.

The result easily follows, for if $x \in f_c(\Gamma_c(r))$, then $x = f_c(z)$ for some $z \in \Gamma_c(r)$ which implies that $x = f_c(z) \in \Gamma_c(r^2)$. Conversely, if $x \in \Gamma_c(r^2)$, then set $z = \sqrt{x-c}$ so that $f_c(z) = x \in \Gamma_c(r^2)$. Then $z \in \Gamma_c(r)$, so $x = f(z) \in f_c(\Gamma_c(r))$, as required. □

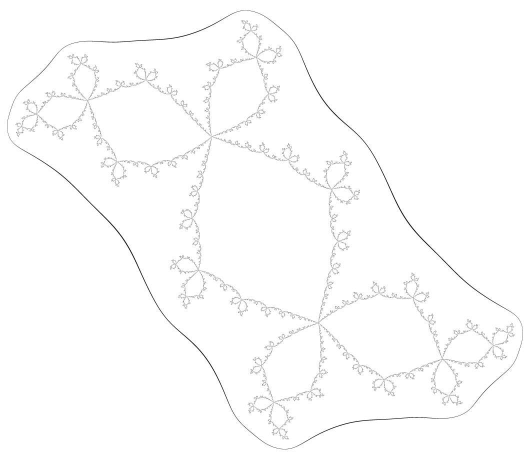

Figure 13: An equipotential of f_ω, wrapped around the Julia set $J(\omega)$. This image was generated using Proposition 3.20 to approximate the value of G_ω for each pixel in the image; if the value is close enough to some constant, we colour the pixel black.

Given the way the equipotential approximates $J(c)$ as $r \to 1$, we may ask what happens exactly in the limit. We will consider this question from two viewpoints. Considering the limit in a local sense, i.e. what happens to a single point on the equipotential curve as $r \to 1$, gives rise to the definition of the dynamical ray. Then, we will change to a more global mindset, and ask if the whole equipotential falls exactly on $J(c)$ in the limit. Through this consideration, we will land smack bang back in the middle of a discussion of local connectivity, due to a theorem of Carathéodory and Torhorst.

So, the dynamical ray is given by fixing the angle t, and varying r. In other words, a dynamical ray is a trajectory that is orthogonal to the equipotentials.

Definition 3.26 Fix $t \in \mathbb{R}/\mathbb{Z}$. The *dynamical ray* of $f_c \in F$, of external angle t, is the set

$$R_c(t) = \{\phi_c^{-1}\left(re^{2\pi it}\right) : r > 1\}$$

A dynamical ray $R_c(t)$ is a line in the complex plane that goes off to infinity, and at its other end gets close to $J(c)$. It is the pre-image under ϕ_c of a straight line going from the unit disc and off to infinity, at an angle of $2\pi t$.

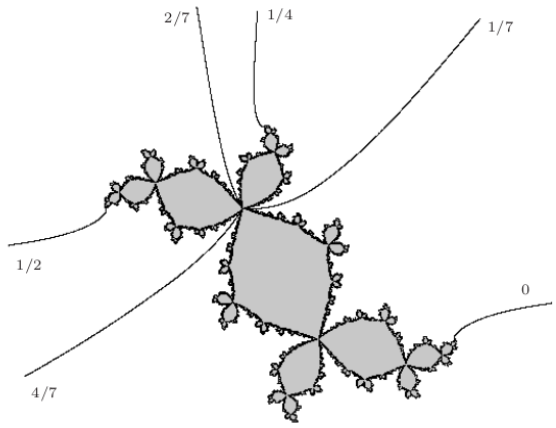

Figure 14: The filled-in Julia set $K(\omega)$, pictured with various dynamical rays labelled by their external angles. This image was taken from page 10 of [Douady and Hubbard, 1985a].

Proposition 3.27 *Let $f_c \in F$. Then $f_c(R_c(t)) = R_c(2t)$.*

Proof. The proof is so entirely in parallel to the proof of Proposition 3.25 that we don't even feel bad about leaving it as an exercise to the reader! (Hint: show $z \in R_c(t) \iff f_c(z) \in R_c(2t)$.) □

Let $R_c(t)$ be some dynamical ray. It is a curve in the complex plane. If we parametrize the ray in the obvious way, we obtain a bijection $(1, \infty) \to R_c(t)$, given by $r \mapsto \phi_c^{-1}\left(re^{2\pi it}\right)$. If the limit

$$z_0 = \lim_{r \to 1} \phi_c^{-1}\left(re^{2\pi it}\right)$$

is well-defined, we say the dynamical ray $R_c(t)$ *lands* at z_0. Note that z_0 is necessarily on the boundary of $K(c)$, that is $z_0 \in J(c)$. We may simply say that a dynamical ray lands, if we know such a point exists but have no need to make explicit reference to it. We say a point $z \in J(c)$ is *accessible* if there is some dynamical ray which lands at it. An accessible point may be the landing point of multiple dynamical rays.

Landing of rays is a central consideration of ours, due to the below version of Carathéodory's Theorem which was first presented and proven in the below form in the thesis of Marie Torhorst [Torhorst, 1921].

Theorem 3.28 *Let $K \subseteq \mathbb{C}$ be compact and connected, with $\mathbb{C}\backslash K$ having no bounded components. Then a conformal isomorphism $\mathbb{C}\backslash K \to \mathbb{C}\backslash D(0,1)$ extends as a homeomorphism to the boundary if and only if K is locally connected.*

When we apply Theorem 3.28 to $K(c)$, we get the following. Note that the local connectivity of $J(c)$ and $K(c)$ is equivalent, because $K(c)\backslash J(c) = \text{int}(K(c))$ is locally connected by a similar argument as in the proof of Theorem 3.13.

Corollary 3.29 *Let $c \in M$. Then $J(c)$ is locally connected if and only if $K(c)$ is locally connected, if and only if all dynamical rays land. Further, if any/all of these conditions hold, the landing points of the dynamical rays $R_c(t)$ vary continuously with t, and are given by the parametrization $\gamma_c = \phi_c^{-1}|_{C(0,1)} : C(0,1) \to J(c)$ of the Julia set of f_c.*

As a result of Corollary 3.29, local connectivity of Julia sets takes on more importance, since it gives a parametrization of the Julia set. This will be put to good use in Section 5.2.

Proposition 3.27 allows us to investigate the dynamics of dynamical rays of Julia sets. In particular, suppose $J(c)$ is locally connected. Then by Proposition 3.27 and since all dynamical rays land, the dynamics of rays under the mapping f_c is topologically conjugate to the angle-doubling map on $C(0,1)$.

It would be desirable to transfer this to the parameter plane, by defining equivalent notions to our equipotentials and dynamical rays. At first, this seems difficult, since the dynamics of f_c were essential in constructing both the explicit form of the Green's function and ϕ_c. However, M is still a compact set in the complex plane, so there exists a Green's function for M. In fact, G_M is in a way induced by the Green's functions G_c, and similarly a conformal isomorphism ϕ_M is induced by the ϕ_c, all of which we make explicit in the following result of Douady and Hubbard [Douady and Hubbard, 1985a]:

Theorem 3.30 *Define two functions, in the parameter plane, by*

$$\begin{aligned} G_M &: \mathbb{C}\backslash M \to \mathbb{R}, & c &\mapsto G_c(c) \\ \phi_M &: \mathbb{C}\backslash M \to \mathbb{C}\backslash D(0,1), & c &\mapsto \phi_c(c) \end{aligned}$$

Then G_M is the Green's function of M, and ϕ_M is a conformal isomorphism.

Recall our observation that c is in the domain of ϕ_c for all $c \in \mathbb{C}$, so that ϕ_M is well-defined. Note also that we have an explicit formula for G_M by Proposition 3.20. Given M's unfathomable complexity, it is remarkable that we can obtain closed forms for these important maps. As Douady and Hubbard put it, "it is rather easier to find the conformal mapping of $\mathbb{C}\backslash M$ than of the complement of a triangle".

Theorem 3.30 allows us to repeat our previous observations in the context of the parameter plane - crucially, including a version of Corollary 3.29 for the parameter plane.

Definition 3.31 Fix $r > 1$. The *parameter equipotential* of M, of potential r, is the set

$$\Gamma_M(r) = \{\phi_M^{-1}\left(re^{2\pi it}\right) : t \in \mathbb{R}/\mathbb{Z}\}$$

Definition 3.32 Fix $t \in \mathbb{R}/\mathbb{Z}$. The *parameter ray* of M, of external angle t, is the set

$$R_M(t) = \{\phi_M^{-1}\left(re^{2\pi it}\right) : r > 1\}$$

A quick note on terminology. While we endeavour in this work to keep the distinction clear, the phrase *external ray* is used to refer to any ray, dynamical or parameter. Similarly, the word *equipotential* is used to refer to an equipotential in either context.

As in the dynamical case, parameter equipotentials are closed simple curves that loop around M, getting tighter as $r \to 1$, while parameter rays are the orthogonal trajectories of the equipotentials. If a parameter ray $R_M(t)$ approaches a well-defined limit $c_0 \in \partial M$ as $r \to 1$, we say the ray *lands* at c_0. A point $c \in \partial M$ is *accessible* if it is the landing point for some parameter ray. An accessible point may be the landing point of multiple parameter rays.

Theorem 3.33 *M is locally connected if and only if ∂M is locally connected, if and only if all parameter rays land. Further, if any/all of these conditions holds, the landing points of the parameter rays $R_M(t)$ vary continuously with t, and are given by the parametrization $\gamma_M = \phi_M^{-1}|_{C(0,1)} : C(0,1) \to \partial M$ of the boundary of M.*

It is known that all parameter rays of rational external angle do indeed land.

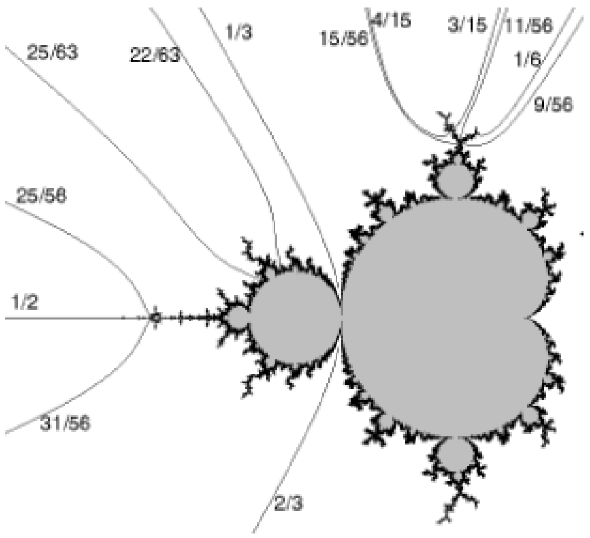

Figure 15: The Mandelbrot set M, pictured with various parameter rays labelled by their external angles. This image was taken from page 2 of [Schleicher, 1997], and was created by John Milnor.

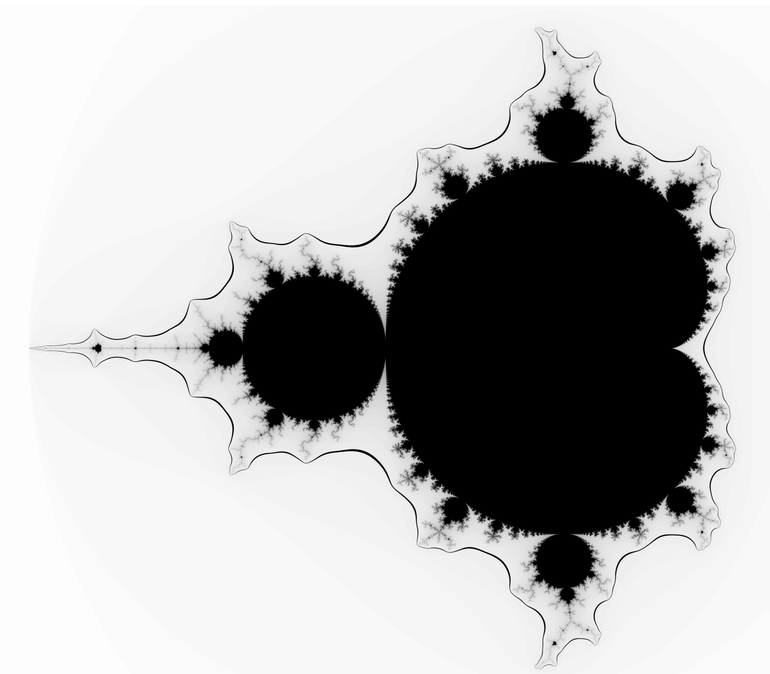

Figure 16: A parameter equipotential, wrapped around M. This image was generated with the same method as Figure 13, utilising Theorem 3.30 and Proposition 3.20 to approximate G_M.

4 The State of the Conjecture Today

Now that we are fully equipped with the necessary background, we are able to delve into the progress of the last few decades. Note that we have already taken our first step in Theorem 3.13, where we proved local connectivity for all parameters inside the Mandelbrot set whose quadratics possess an attracting cycle.

Over the past 36 years, there have been tremendous efforts to whittle down the unsolved cases of the conjecture. Grand mathematical machinery has been crafted to overcome wide expanses of difficulties, with a resolution to the conjecture being provided for more and more classes of parameters. We take a journey through time, recording a comprehensive timeline of results towards the MLC conjecture. Afterwards, we focus in on (a very narrow subset of) the details of the related proofs.

4.1 Timeline of the Conjecture

1984 [Douady and Hubbard, 1985a]
Douady and Hubbard assume the Mandelbrot set is locally connected to prove density of hyperbolicity, creating our eponymous conjecture and firing the starting pistol on a monumental body of work for the following decades in holomorphic dynamics.

1985 [Douady and Hubbard, 1985b]
The concept of polynomial-like mappings is introduced by Douady and Hubbard to explain the self-similarity and universality of the Mandelbrot set. This would become intimately linked with the MLC conjecture via quadratic-like renormalization (see Section 4.3), due to the below result.

1989 [Hubbard, 1993]
With an enormous breakthrough, Yoccoz proves MLC at all parameters that are either: (*i*) not ql-renormalizable, (*ii*) finitely ql-renormalizable, (*iii*) Misiurewicz points or (*iv*) on the boundary of a hyperbolic component. The results are published a few years later by Hubbard. As a result, it remains only to show local connectivity at parameters of quadratics that are infinitely ql-renormalizable.

1993 [Lyubich, 1997]
For the first time, due to Lyubich, we see parameters of infinitely ql-renormalizable quadratic maps of bounded type for which MLC holds. In particular, he shows local connectivity at all parameters of infinitely ql-renormalizable quadratic maps that satisfy the following two conditions:

- (i) The hybrid classes of all the renormalizations of the map are picked from a finite number of truncated secondary limbs, and
- (ii) The combinatorial type of the map is sufficiently high on all levels, depending on the limbs chosen in condition (i).

1995 [Jiang, 1995]
Jiang constructs a subset of parameters whose quadratics are infinitely ql-renormalizable and dense on the boundary of M, for which MLC holds.

1999 [Schleicher, 1999]
Schleicher gives new proofs of local connectivity at Misiurewicz parameters and for parameters on the boundary of a hyperbolic component, by introducing fibers of the Mandelbrot set.

2006 [Kahn, 2006]
Kahn proves that MLC holds at all paramaters whose maps are infinitely primitively ql-renormalizable of bounded type.

2007 [Kahn and Lyubich, 2006]
Kahn and Lyubich demonstrate local connectivity at all parameters whose quadratics are infinitely ql-renormalizable of primitive type and satisfy the *decoration* condition, which means the combinatorics of the renormalization operators involved is selected from a finite family of decorations, which are parts of M that are cut off from the rest of the set by a pair of parameter rays which land at the tip of a satellite copy of M.

2007 [Levin, 2007]
Levin publishes proofs of local connectivity for parameters whose quadratics are infinitely ql-renormalizable with the associated rotation numbers of separating fixed points of the renormalizations obeying a certain inequality in an upper limit.

2008 [Kahn and Lyubich, 2008]
Kahn and Lyubich extend their work to give MLC at all parameters whose quadratics are infinitely ql-renormalizable of primitive type and satisfy the *molecule* condition, which means the combinatorics of the primitive renormalization operators involved stays away from the molecule, which is the part of M comprising of the least full superset of the closure of the main cardioid and all hyperbolic components that can be reached from it via bifurcation points. The molecule condition means that all renormalizations stay uniformly away from the satellite type.

2011 [Levin, 2011]
Levin extends his work, relaxing the conditions of [Levin, 2007] to give a larger class of parameters for which MLC holds.

2015 [Cheraghi and Shishikura, 2015]
Cheraghi and Shishikura use the *Almost Parabolic Renormalization* of Inou and Shishikura [Inou and Shishikura, 2008] to prove MLC for parameters whose quadratics are infinitely ql-renormalizable of unbounded satellite type, given certain growth conditions on the combinatorics.

2019 [Dudko and Lyubich, 2019]
Using *Pacman Renormalization*, Dudko and Lyubich publish the latest breakthrough in the conjecture and show that M is locally connected at certain parameters whose quadratics are infinitely ql-renormalizable of bounded satellite type. This is the first example of local connectivity at such parameters.

4.2 Quadratic-like Renormalization

The above timeline gives a strong indication of the relevance of *renormalization* to the MLC conjecture. This link was first established by Yoccoz, who proved local connectivity of M for all points other than the *infinitely ql-renormalizable* ones. We will outline this result, but first it is of course necessary to investigate the quadratic-like renormalization of Douady and Hubbard. This itself depends on maps between simply connected domains. Thus, we feel the need to develop a more rigorous footing here than that provided in Definition 2.19. This brings us back to a discussion of *conformal mappings*, which is probably also in need of a slightly more rigorous treatment than the loose definition given earlier, especially given the importance of Theorem 3.15.

Conformal mappings

Recall from Definition 3.14 that a *conformal mapping* is a mapping between two subsets D, E of the complex plane that preserves angles, and a *conformal isomorphism* is a conformal mapping that is also a bijection. We have as yet avoided laying this concept out in a rigorous way, since there is a neater, equivalent way of thinking about conformal mappings.

Before getting into this section, we warn the reader that various authors define conformal mappings in different, and not equivalent, ways. We avoid the gory details, and will use the below result as our standard of a conformal mapping, while keeping in mind the angle-preserving nature of these maps.

Theorem 4.1 (Conformal Mapping Theorem) *Let $A, B \subseteq \mathbb{C}$ be open. Then $f : A \to B$ is a conformal mapping if and only if f is holomorphic with $f'(z) \neq 0$ for all $z \in A$.*

To add to the confusion, we have an alternative characterisation of conformal isomorphisms. Note the below results very much depend on our choice of definition for a conformal mapping. We include these results here in an attempt to avoid ambiguity for the keen reader who intends to read this subject more widely. We like this reader!

Definition 4.2 Let $A, B \subseteq \mathbb{C}$ be open. Then $f : A \to B$ is *biholomorphic* if it is a bijection, with both f and f^{-1} being holomorphic.

Note that f is biholomorphic if and only if f^{-1} is biholomorphic.

With our definitions, the Inverse Function Theorem for Holomorphic Maps states that conformal isomorphisms are biholomorphic. Further, since it may be shown that injective complex maps defined on open subsets of $\mathbb{C}$ must have everywhere non-zero derivative, the converse also holds.

Proposition 4.3 *Let $A, B \subseteq \mathbb{C}$ be open. Then $f : A \to B$ is biholomorphic if and only if it is a conformal isomorphism.*

So, to summarise, there is no standard definition of conformal mappings on which all sources agree, and we will take conformal mappings to be the holomorphic maps with non-zero derivative. If we then define a conformal isomorphism to be a bijective conformal mapping, then conformal isomorphisms and biholomorphisms are exactly the same thing.

Conformal isomorphisms give a classification of subsets of the complex plane.

Definition 4.4 If there is a conformal isomorphism $f : A \to B$, we say A and B are conformally equivalent.

Proposition 4.5 *Conformal equivalence is an equivalence relation.*

Proof. Let $A, B, C \subseteq \mathbb{C}$ be open subsets.

Reflexivity: The identity map $A \to A$ given by $z \mapsto z$ is clearly bijective, and holomorphic with non-zero derivative. Hence by Theorem 4.1, it is a conformal isomorphism.

Symmetry: Suppose there is a conformal isomorphism $f : A \to B$. By Proposition 4.3, f is biholomorphic. Thus by the definition of a biholomorphic map, $f^{-1} : B \to A$ is also biholomorphic, and so again by Proposition 4.3, this is a conformal isomorphism $B \to A$.

Transitivity: Suppose there are conformal isomorphisms $f : A \to B$ and $g : B \to C$. Compositions of bijective holomorphisms are themselves bijective holomorphisms, so that $g \cdot f : A \to C$ is a bijective holomorphism. By the chain rule, $(g \circ f)'(z) = g'(f(z)) \cdot f'(z)$. Since f' and g' are always non-zero, we see that $(g \circ f)'$ is also. Thus, $(g \circ f)$ is a conformal isomorphism $A \to C$. □

For many purposes, we will think of conformally equivalent subsets of the plane as 'being the same' as each other. It is a stronger form of equivalence than homeomorphicity. We will be especially interested in those subsets of $\mathbb{C}$ which are conformally equivalent to $B(0, 1)$, which brings us to the most important theorem of this subsection.

Theorem 4.6 (Riemann Mapping Theorem) *Let $U \subset \mathbb{C}$ be open and simply connected. Then U is conformally equivalent to the open unit disc $B(0, 1)$.*

Note that we require U to be a strict subset of the plane: Liouville's Theorem may be used to show that $\mathbb{C}$ and $B(0, 1)$ are not conformally equivalent.[8] It follows from Theorem 4.6 that all strict subsets of $\mathbb{C}$ which are open and simply connected are conformally equivalent.

The below corollary will be how we deal with simple connectivity, in place of the imprecise definition we gave back in Section 2.3.

Corollary 4.7 *Let $U \subseteq \mathbb{C}$ be open. Then U is simply connected if and only if $U = \mathbb{C}$ or U is conformally equivalent to $B(0, 1)$.*

[8]It so happens that this already brings us very close to a complete classification of simply connected 'Riemann surfaces', up to conformal equivalence. For more on this, see the *Uniformization Theorem* (Theorem 1.1 of [Milnor, 1990]).

Proof. (⇒) Restatement of Theorem 4.6.
(⇐) $\mathbb{C}$ and $B(0,1)$ may both be seen to be simply connected. Then, suppose then U is conformally equivalent to $B(0,1)$ via, say $\phi : B(0,1) \to U$. Since ϕ is a homeomorphism, and since it may be shown that simple connectivity is a topological property (i.e. invariant under homeomorphisms), U must also be simply connected. □

In particular, when we talk about open subsets of $\mathbb{C}$ which are conformally equivalent to $B(0,1)$, as we frequently will, we are talking precisely about those open proper subsets of $\mathbb{C}$ which are simply connected.

Quadratic-like renormalization: the self-similarity, and universality, of M

To overcome the final barrier to our locally-connected endeavours, we introduce the idea of quadratic-like mappings. This will equip us with the power of quadratic-like renormalization, which can be used to explain the presence of small, approximate copies of the Mandelbrot set within itself. In the course of our discussion, we will touch on the universality of the Mandelbrot set - this is one of the biggest reasons for the enduring mathematical importance of M, and therefore lends importance to the MLC conjecture.

Renormalization is a sort of catch-all term for a very general process, of which there is a large variety of specific forms. The general idea is that dynamical systems often exhibit a repetition of form within themselves. For example, say we are investigating the dynamics of some $f : X \to X$ on some space X. There may be some small part of the space $U \subset X$ which iterates of f return to periodically, say $f^{\circ k}(U) = U$. Then defining $g = f^{\circ k}|_U$, we have dynamics of g, in a sense within f's dynamics and on a smaller scale. In general, the dynamics of such a g will differ from the dynamics of f (though not always– see Example 4.16). However, if f originated from some family of functions, it might be that g has dynamics resembling those of another member of the family. The 'renormalization step' is to find such a map in the family. Renormalization is the process of going from a map, to some scaled down dynamics within, and then scaling back up by passing to the map found in the renormalization step.

We should note quickly that there are many different renormalizations defined for families of holomorphic functions, and even just in the world of the MLC conjecture (see various results in Section 4.1). Thus, unless it is otherwise stated, the reader should assume any use of the word renormalization refers to the quadratic-like renormalization.

The quadratic-like renormalization process for complex quadratic maps is as follows. We start with some $f_c \in F$, which has to have one of its iterates satisfy certain properties when restricted to an appropriate subset of $\mathbb{C}$. This restricted function will then have some of the properties of a complex quadratic map ('quadratic-like'), and the renormalization step is to apply the *Douady-Hubbard Straightening Theorem* which says there is another $f_{c'} \in F$ which shares the dynamics of the restricted function.

Although we are defining quadratic-like mappings because of its links via quadratic-like renormalization to results relating to the MLC conjecture, their most well-known application is to explain the phenomenon of approximate copies of the Mandelbrot set appearing in parameter spaces for other function spaces, as well as in itself (Figure 3). This was Douady and Hubbard's motivation for introducing quadratic-like mappings in [Douady and Hubbard, 1985b]. This will be briefly visited en route.

For our discussion, we will require the following technical notion.

Definition 4.8 Let $A, B \subseteq \mathbb{C}$, and let $f : A \to B$. Then f is *proper* if, for each compact subset $K \subseteq B$, we have that the pre-image $f^{-1}(K)$ is compact in A.

It will be crucial to us that all (non-trivial) iterates of complex quadratics are proper:

Proposition 4.9 *Let $A, B \subseteq \mathbb{C}$. Suppose $f : A \to B$ is a non-constant polynomial, say $f(z) = \sum_{k=0}^{n} a_k z^k$ with $n \geq 1$ and $a_0, \ldots, a_n \in \mathbb{C}$. Then f is a proper map. In particular, any restriction of $f_c^{\circ k}$ is a proper map for all $c \in \mathbb{C}, k \geq 1$.*

Proof. We begin by noting that f is indeed holomorphic. Let $K \subseteq B$ be compact. Then as a compact subset of the complex plane, K is closed and bounded. Consider $f^{-1}(K)$. Since f is a polynomial, it is holomorphic and in particular continuous. Therefore, f is a closed map, so that $f^{-1}(K)$ is closed. Next, we have $K \subseteq B(0, r) \cap A$ for some $r > 0$, since K is bounded. Therefore, $f^{-1}(K) \subseteq f^{-1}(B(0, r))$. If A is bounded, then so too is $f^{-1}(K) \subseteq A$. If A is unbounded, suppose for a contradiction that $f^{-1}(K)$ is unbounded. Then we may choose points in $f^{-1}(K)$ with arbitrarily large absolute value. Since $|f(z)| \to \infty$ as $|z| \to \infty$, we may choose some $z \in f^{-1}(K)$ such that $|f(z)| > r$. But then $z \notin f^{-1}(B(0, r))$, contradicting $f^{-1}(K) \subseteq f^{-1}(B(0, r))$. We conclude that $f^{-1}(K)$ is closed and bounded, so that it is compact, and therefore f is a proper map. $\square$

Examples of non-proper maps include the constant functions $z \mapsto z_0$ for $z_0 \in \mathbb{C}$ (since the pre-image of the compact set $\{z_0\}$ is the unbounded set $\mathbb{C}$) and the trigonometric function $\sin : \mathbb{C} \to \mathbb{C}$ (since the pre-image of the compact set $\{0\}$ is the unbounded set $\{n\pi : n \in \mathbb{Z}\}$). In fact, it may be shown that the non-constant polynomials are the only holomorphic proper maps $\mathbb{C} \to \mathbb{C}$.

The importance of proper maps to us, is the idea of its *degree*. For a proper map $f : U \to V$, every point in V that is not a critical value of f has the same number of pre-images. This fixed number of pre-images is the degree. We will not get into this any more deeply, let alone prove this property of proper maps. Instead, we will define quadratic-like maps with the idea of multiplicity of pre-images of polynomials, which loses some generality but suffices for investigating quadratic-like renormalization. The key takeaway is that, if we can show a holomorphic map to be proper, and that a non-critical value has 2 pre-images, then the below definition is satisfied.

Definition 4.10 Let U, V be open subsets of $\mathbb{C}$ that are conformally equivalent to $B(0, 1)$, which are such that $\overline{U} \subset V$. A *quadratic-like mapping* is a holomorphic proper map $f : U \to V$ such that every point in V has exactly two pre-images in U (counted with multiplicity).

Recall that the multiplicity of a root α of a polynomial $P(z)$ over $\mathbb{C}$ is given as the highest power of $(z - \alpha)$ that divides $P(z)$. There is a general way to define the multiplicity of a pre-image of a complex map, however the only quadratic-like maps we will only ever be dealing with are (suitable restrictions of) iterates of quadratics - that is, polynomials of degree 2^k for some $k \in \mathbb{N}$ - and so we may define the multiplicity of a pre-image in the following simpler way: for a polynomial $f : U \to V$, suppose $u \in U$ is a pre-image of $v \in V$, so that $f(u) = v$. Then the multiplicity of u as a pre-image of v is the multiplicity of u as a root of the polynomial $f(z) - v$. Counting multiplicities of pre-images in this way, our definition of quadratic-like maps coincides with requiring f be a holomorphic proper map of degree 2. As a final note on the definition, note that a quadratic-like map f is necessarily surjective.

We are very keen on looking at some examples of quadratic-like maps, as it feels like a lot to take in. We choose to do this after stating Theorem 4.12, as this theorem gives such a huge insight into such examples. One example for which that theorem gives no new information is that quadratic maps themselves are quadratic-like, so we look at this example now.

Example 4.11 The 'two pre-images' property of quadratic-like maps is satisfied by quadratic maps, given we define the quadratic on the right domain. Accordingly, suitable restrictions of quadratic maps are quadratic-like. Let $f_c \in F$, and define U to be the unique open set whose boundary is the equipotential $\Gamma_c(r)$ for some $r > 1$ and which contains $J(c)$. Similarly, take V to be the unique open set whose boundary is the equipotential $\Gamma_c(r^2)$ and which contains $J(c)$, and consider the restricted map $f_c|_U : U \to V$. We claim this is a quadratic-like map. By Proposition 3.25, it does indeed map into V, and it may be shown these sets are conformally equivalent to $B(0, 1)$, and that $\overline{U} \subset V$. By Proposition 4.9, it is holomorphic and proper. Let $z \in K(c) \backslash \{c\}$. Then since f_c is a quadratic, z has two pre-images which by Proposition 2.10 are in $K(c) \subseteq U$. Thus, $f_c|_U$ is a quadratic-like map.

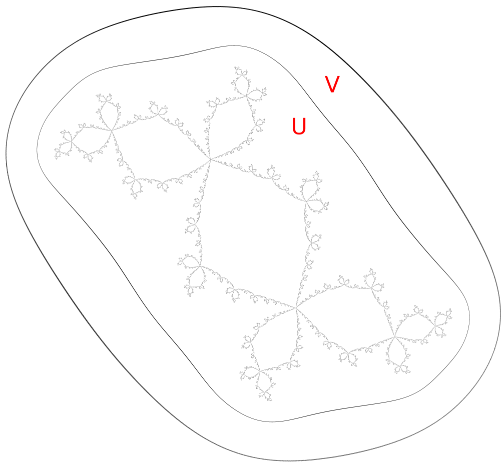

Figure 17: Consider the Julia set $J(\omega)$, as in Example 2.2 and Figure 4). It is pictured, along with two curves. The inner curve is the equipotential $\Gamma_\omega(1.2)$, and the outer curve is its image $f_\omega(\Gamma_\omega(1.2)) = \Gamma_\omega(1.44)$. Taking U as the interior of the region bounded by $\Gamma_\omega(1.2)$, and V as the interior of the region bounded by $\Gamma_\omega(1.44)$, the map $f_\omega|_U : U \to V$ is a quadratic-like map.

Just as with quadratic maps, we may investigate the dynamics of quadratic-like maps. Many of the definitions and theorems for the dynamics of quadratic maps carry across to the more general class of quadratic-like maps, often with effectively identical proofs. In particular, we may define the filled-in Julia set of a quadratic-like map $f : U \to V$ as

$$K(f) = \{z \in U : f^{\circ k}(z) \in U \text{ for all } k \in \mathbb{N}\}$$

and the Julia set of f as

$$J(f) = \partial K(f).$$

Note the subtle change in both definition, and notation. Previously, the filled-in Julia set consisted of all points which don't go off to infinity, whereas now we require they remain strictly in U, since otherwise our map f is not defined. Also, since quadratic-like maps have a far more general form than quadratic maps, there is no one-variable parametrisation, so the notation $K(f)$ specifies the actual name of the quadratic-like map rather than the $K(c)$ we have been using for $f_c \in F$.

Those quadratic-like mappings with connected Julia set will be of particular interest. Just as in the quadratic case, the Julia set of a quadratic-like map is connected if and only if every critical point of the map belongs to its filled-in Julia set.

We now introduce the most important theorem regarding quadratic-like mappings, the so-called *Straightening Theorem* of Douady and Hubbard. It is essential to the process of quadratic-like renormalization, and also explains why the above definitions and results for quadratic-like maps are so in parallel with those for quadratic maps.

Theorem 4.12 (Straightening Theorem for Quadratic-like Mappings) *Let U, V be open subsets of $\mathbb{C}$ that are conformally equivalent to $B(0,1)$, and which satisfy $\overline{U} \subset V$, and let $f : U \to V$ be a quadratic-like mapping. Then f is hybrid equivalent to some quadratic map. Moreover, if $J(f)$ is connected, then this quadratic map is unique up to affine conjugation. In particular, in the connected case, f is hybrid equivalent to a unique $f_c \in F$.*

Hybrid equivalence is a stronger form of topological conjugacy. In particular, if f and g are hybrid equivalent, then there is a homeomorphism $\phi : A \to B$, where A and B are neighbourhoods of $K(f)$ and $K(g)$ respectively, such that $\phi \circ f = g \circ \phi$, and ϕ is *quasi-conformal* (doesn't distort angles too much) with the *Wirtinger derivative* $\frac{1}{2}\left(\frac{\partial \phi}{\partial x} + i\frac{\partial \phi}{\partial y}\right)$ evaluating to 0 throughout $K(f)$. We avoid the details - the crucial point is that hybrid equivalent maps are topologically conjugate. As a result, they have essentially identical dynamics, and homeomorphic Julia sets.

The following example is illustrative of the self-referential nature of M, and the power of the Straightening Theorem in explaining this phenomenon.

Example 4.13 Let $c \approx -1.7577 + 0.0134i$, so that $f_c^{\circ 9}(0) = 0$. Certainly $c \in M$.

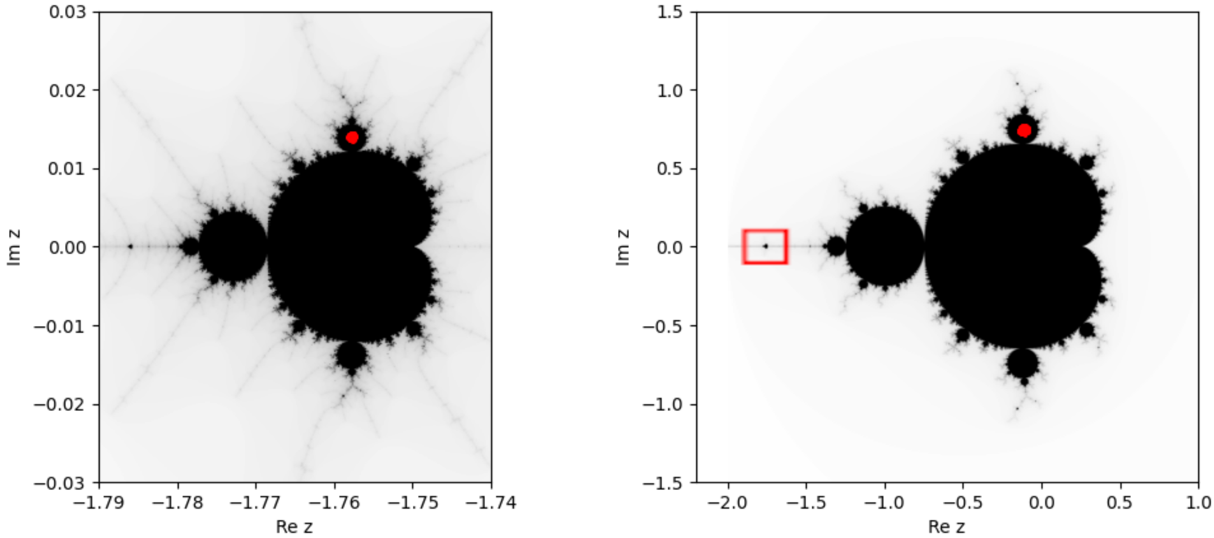

Figure 18: Left: We zoom in on a homeomorphic copy of the Mandelbrot set within itself. The approximate location of the parameter c for this example is shown as a red dot.

Right: The entirety of the Mandelbrot set. The 'mini Mandelbrot' shown on the left is just about visible within the red box. The approximate location of the parameter ω is shown as a red dot.

The Julia set of c is particularly striking:

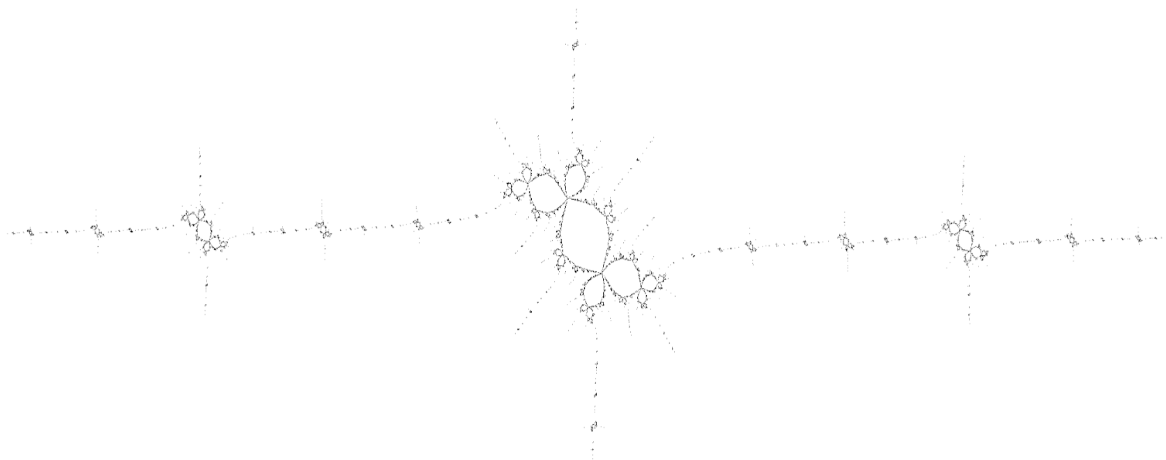

Figure 19: The Julia set $J(c)$. It is seemingly made up of scaled-down Douady rabbits, joined by thin filaments resembling those on the mini-Mandelbrot.

If we plot the critical orbit $O(c, 0)$ on top of the Julia set, the intrigue only deepens...

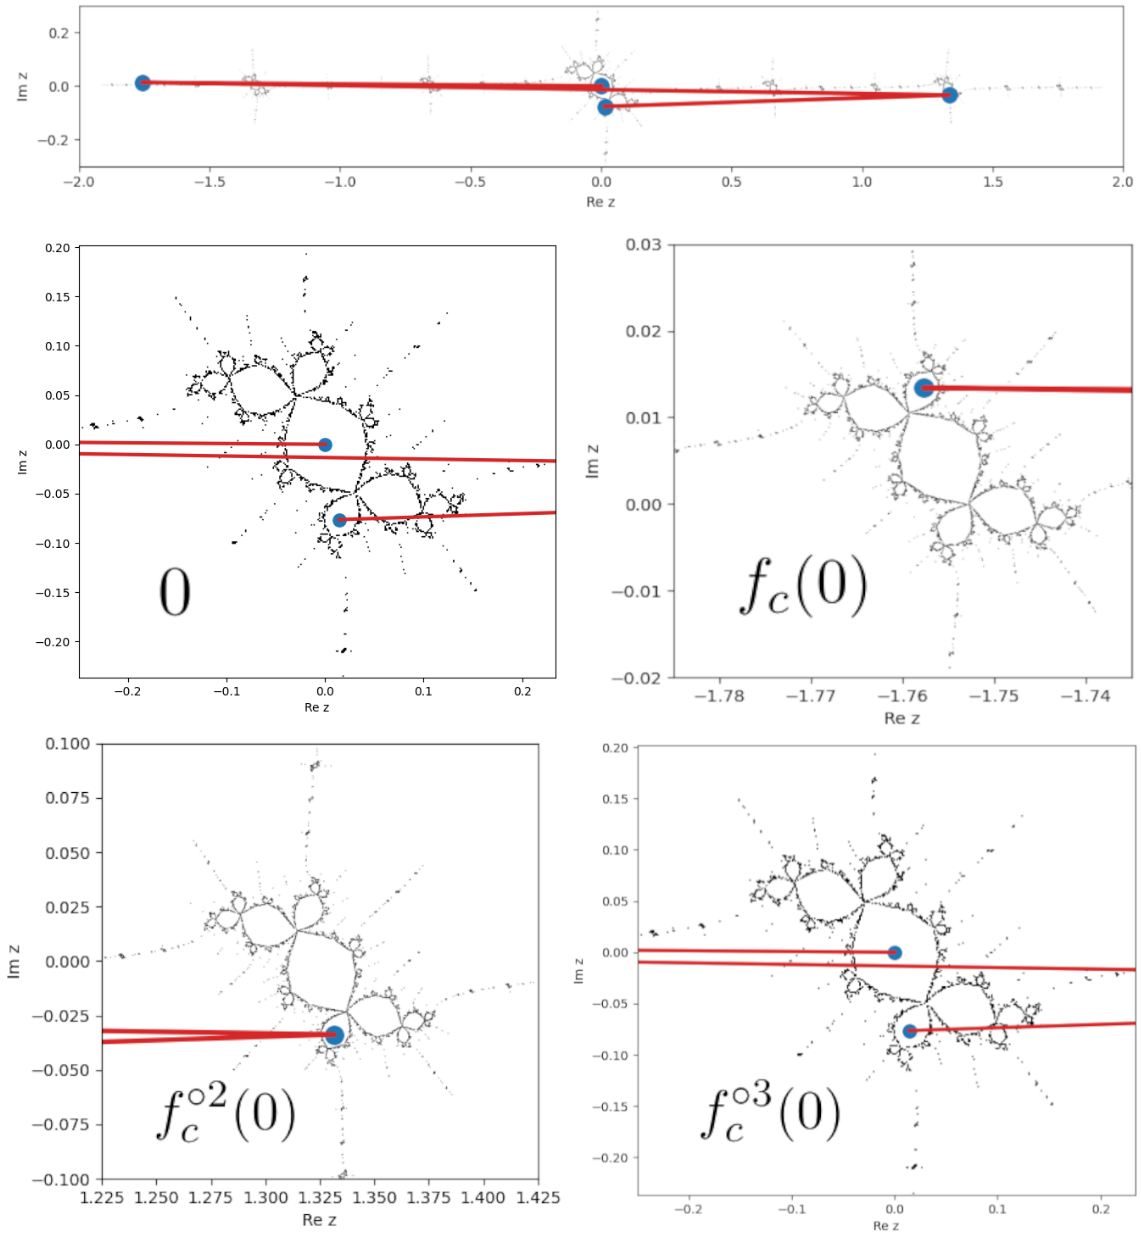

Figure 20: We sequentially show the first 4 points in the critical orbit of f_c, on its Julia set. In this way, we can see that $f_c^{\circ 3}$ is the first 'return' map for the central mini-rabbit.

The orbit jumps between 3 particular 'mini-rabbits', and this pattern will hold for the entire orbit, and for any starting point that is within one of the 3 rabbits. The central rabbit, where we start and which contains the critical point 0, is of particular interest. It is contained within the open set

$$U = \{z \in \mathbb{C} : -0.3 < \mathrm{Re}(z) < 0.3 \text{ and } -0.3 < \mathrm{Im}(z) < 0.3\}$$

If we consider $f_c^{\circ 3}$ as a function, we will 'skip over' the other mini-rabbits in the dynamics, and jump around this central rabbit, within U. Letting $V = f_c^{\circ 3}(U)$ and $f = f_c^{\circ 3}|_U : U \to V$, we are led to the conclusion that f is a quadratic-like map. We may solve the equation $\left(f_c^{\circ 3}\right)'(z) = 0$ numerically to find that $f_c^{\circ 3}$ has only one critical point in U, namely 0. Thus, 0 is the unique critical point of f. Since $f^{\circ 3}(0) = f_c^{\circ 9}(0) = 0$, we must have $0 \in K(f)$, so that $J(f)$ is connected, and by Theorem 4.12, f is hybrid equivalent to a unique $f_c \in F$. In this case, f is hybrid equivalent to the Douady rabbit quadratic f_ω. Thus, in this case, $J(f)$ is a small homeomorphic copy of the Douady rabbit $J(\omega)$.

If we let c vary over the mini-Mandelbrot that we have been considering, we would be varying over a collection of quadratic-like maps. Each one would be hybrid equivalent to a unique $f_{c'} \in F$. The map

$c \mapsto c'$ may be shown to be a homeomorphism, so that we are indeed dealing with a homeomorphic copy of M. This process can be applied to any family of quadratic-like maps. This leads to homeomorphic copies of M appearing all over the place in one-dimensional complex dynamics. This is the so-called *universality* of the Mandelbrot set.

Definition 4.14 Let $f_c \in F$. Then f_c is *renormalizable at level n* if there exist open sets $0 \in U, V \subset \mathbb{C}$ that are conformally equivalent to $B(0,1)$, such that for some $n \in \mathbb{N}$, we have a quadratic-like map $f_c^{\circ n} : U \to V$ which has a connected Julia set. The unique $f_{c'} \in F$ which is hybrid equivalent to the quadratic-like map $f_c^{\circ n}$ is a *quadratic-like renormalization* of f_c. We denote this by $\mathcal{R}_n(f_c) = f_{c'}$, so that $\mathcal{R}_n$ is the *quadratic-like level-n renormalization operator* which maps level-n renormalizable quadratics to their renormalization in F. If f_c is renormalizable at level n for some $n \in \mathbb{N}$, we say f_c is *renormalizable.* Otherwise, f_c is *not renormalizable.*

Looking back at Example 4.13, with c as in that example, we showed that f_c is renormalizable, with renormalization $\mathcal{R}_3(f_c) = f_\omega$.

Suppose we take some renormalizable $f_c \in F$ to its renormalization $\mathcal{R}_n(f_c)$. What if the renormalization itself is again renormalizable? We are motivated to make a definition!

Definition 4.15 Let $f_c \in$ be renormalizable at level n. If $\mathcal{R}_n(f_c)$ is renormalizable, with a renormalization which is again renormalizable, and so on, then we say f_c is *k-times renormalizable*, where $k \in \mathbb{N}$ is the number of times we are able to repeat the renormalization process. If f_c is k-times renormalizable for infinitely many $k \in \mathbb{N}$, we say f_c is *infinitely renormalizable.* Otherwise, f_c is *not infinitely renormalizable.* If f_c is k-times renormalizable for some $k \in \mathbb{N}$, but is not infinitely renormalizable, then we say f_c is *strictly finitely renormalizable.*

This definition of *infinitely renormalizable* is equivalent to the statement that $f_c \in F$ is renormalizable at level n for infinitely many $n \in \mathbb{N}$.

The mini-Mandelbrot considered in Example 4.13 has smaller Mandelbrot copies attached to it, each of which has yet-smaller Mandelbrot copies, and so on. A parameter that was located in a hyperbolic component of a 'k-level' mini-Mandelbrot would be k-times renormalizable, with each renormalization pulling us out to the next biggest mini-Mandelbrot.

An infinitely renormalizable map has an infinity of layers to its dynamics. A parameter that was within infinitely many mini-Mandelbrots would give an infinitely renormalizable map.

Example 4.16 Let $c_F = -1.4011551890\ldots$ be the so-called *Feigenbaum parameter.* The location of the parameter is easily seen on a picture of M. Starting at 0 at the centre of the main cardioid, head left along the real axis. The first 'period doubling bifurcation parameter' is $c = -\frac{3}{4}$, and is where the main cardioid meet the unique hyperbolic component of period 2. It is so-called because crossing it takes us from quadratics with attracting cycles of period 1 to quadratics with attracting cycles of period 2. Continuing left along the real axis, we will cross the next period doubling bifurcation point when we cross over to the next hyperbolic component, which is of period 4, and going on this way, we will find infinitely many period doubling bifurcations in a compact interval. The Feigenbaum parameter c_F is defined as the smallest negative real number such that there are infinitely many period doubling bifurcation points in $[c_F, 0]$. It is the limit point of the set of period doubling bifurcation parameters. Let c_n be the n^{th} period doubling bifurcation parameter. Then c_n is the root of a hyperbolic component W_n of period 2^n, which is a perfect circle. The ratio diam W_k/diam W_{k+1} of the diameters approaches the famous *Feigenbaum constant* $\delta = 4.6692016\ldots$

f_{c_F} may be shown to be renormalizable, with $\mathcal{R}_2(f_{c_F}) = f_{c_F}$. That is, it renormalizes to itself. Equivalently, we have that $f_{c_F}^{\circ(2^k)}$ is quadratic-like when restricted to a suitable domain, for every $k \in \mathbb{N}$. It follows that f_{c_F} is infinitely renormalizable.

It is known the Julia set $J(c_F)$ is locally connected, but it is at this moment unknown whether M is locally connected at c_F (see [Hu and Jiang, 1993]).

This map remains a subject of research - as recently as 2020, Artem Dudko and Scott Sutherland published their result that the Hausdorff dimension of $J(c_F)$ is less than 2. This had previously been a long-standing open problem (see [Dudko and Sutherland, 2020]).

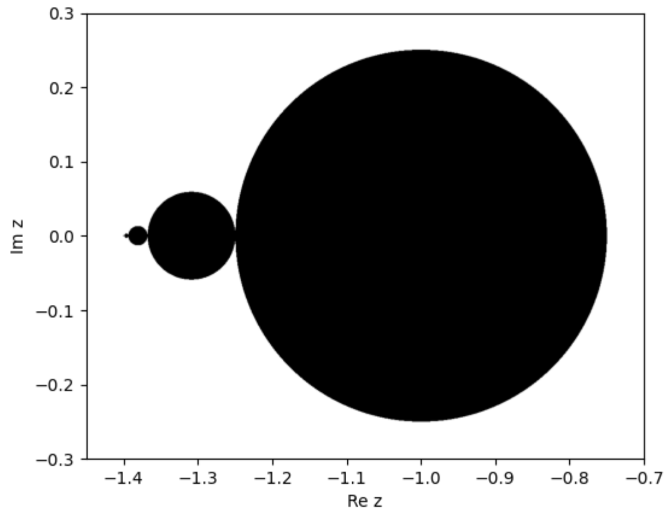

Figure 21: The circular hyperbolic components W_n of M, which are each the unique hyperbolic component of period 2^n that intersects the real line. Here we have drawn W_n for $1 \leq n \leq 9$. The ratios of the diameters is tending to the Feigenbaum constant δ. The left-most point of W_n tends to the Feigenbaum parameter c_F.

4.3 Yoccoz' Theorem

Here, we outline the proofs as presented by Hubbard in [Hubbard, 1993] of J. C. Yoccoz' breakthrough results of 1989. By the end of the section, we will have local connectivity at all non-infinitely quadratic-like renormalizable parameters.

Theorem 4.17 (Yoccoz' Theorem) *Suppose $f_c \in F$ is not infinitely renormalizable. Then M is locally connected at c.*

It is a long journey to get to this result. The path there is as follows: First, we show that M is locally connected at all parameters whose quadratics possess an indifferent cycle. Then, we are free to focus only on quadratics whose cycles are all repelling. If this is the case for some f_c which is not infinitely renormalizable, then the Julia set $J(c)$ is locally connected. The final step is to transfer this result to the parameter plane, giving that M is locally connected at the corresponding c.

The 'indifferent parameters'

An *indifferent parameter* is a parameter $c \in M$ such that f_c possesses a cycle whose multiplier ρ satisfies $|\rho| = 1$. In spite of the name, there are many interesting things to be said of these parameters. For example, the following generalisation of Corollary 3.11. The idea behind the proof is to suppose there are two non-repelling cycles, and show that we may perturb the parameter so that both become attracting, contradicting Corollary 3.11.

Proposition 4.18 (Fatou-Shishikura Theorem for quadratic maps) *Let $f_c \in F$. Then f_c has at most one cycle which is not repelling.*

Thus, for an indifferent parameter, we may talk of *the* indifferent cycle that belongs to it. Indifferent parameters necessarily belong to the boundary of M. In fact, if f_c has an indifferent cycle of length

k, then c belongs to the boundary of a hyperbolic component W which consists of parameters with attracting cycles of length k.

The situation is particularly clear when the cycle is rationally indifferent. In this case, we have exactly two parameter rays landing at the indifferent parameter - with the exception of the 'cusp' parameter $c = \frac{1}{4}$, which is the landing point of $R_M(0)$ only. Excluding this point from our considerations, we are able to use this fact about rationally indifferent parameters to construct *wakes* and *limbs* of the Mandelbrot set. These will be used to construct arbitrarily small connected neighbourhoods for (most) indifferent parameters.

Definition 4.19 Suppose $f_c \in F$ has a rationally indifferent cycle with multiplier $\rho = e^{2\pi i p/q}$, with $0 < p < q$ co-prime. Let $R_M(t_1)$ and $R_M(t_2)$ be the two parameter rays landing at c. Consider the set

$$\mathbb{C} \setminus (R_M(t_1) \cup R_M(t_2) \cup \{c\})$$

constructed by taking the complex plane, and removing both the ray pair and their landing point. It has two connected components. The *wake* of c is the connected component which does not contain 0.

Example 4.20 Consider the first period-doubling bifurcation parameter $c = -\frac{3}{4}$. It is the unique point in the intersection of the closures of the unique hyperbolic components of periods 1 and 2. As such, it is an indifferent parameter. Its unique indifferent cycle is the fixed point $z = -\frac{1}{2}$. The parameter is the landing point of the two parameter rays $R_M(1/3)$ and $R_M(2/3)$. This is all shown in Figure 22, where the resulting wake is the part of the plane on the left of the rays.

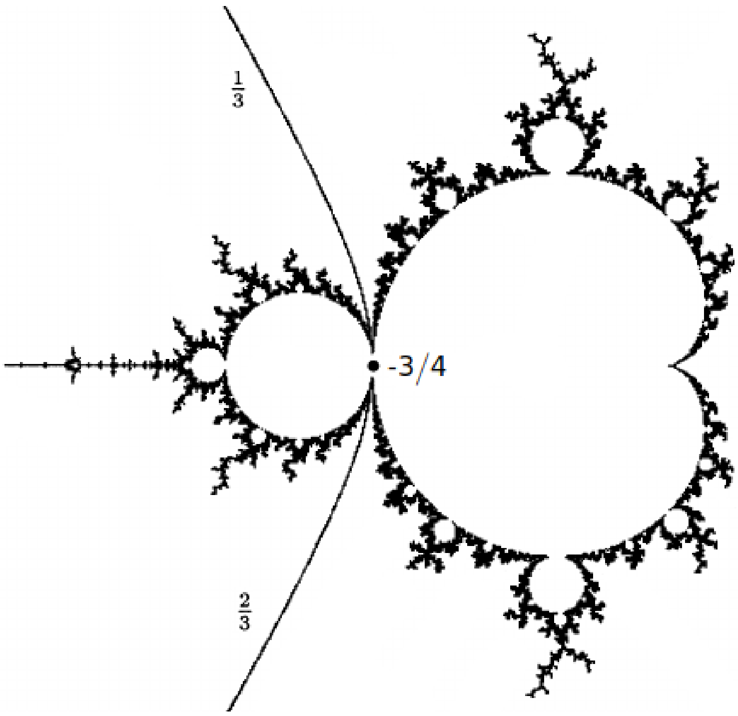

Figure 22: This image shows the Mandelbrot set along with the parameter rays $R_M(1/3)$ and $R_M(2/3)$, which land at the indifferent parameter $c = -\frac{3}{4}$. The ray pair cuts the parameter plane into two connected components. The component not containing zero, that is the one disjoint from the main cardioid, is the wake of c. This image has been modified from the form it takes in its source, which is [Branner and Fagella, 2001] - there, it is Figure 1 on page 101.

In case you didn't quite get the idea from Example 4.20, let us take a brief look at roughly 1000 more examples!

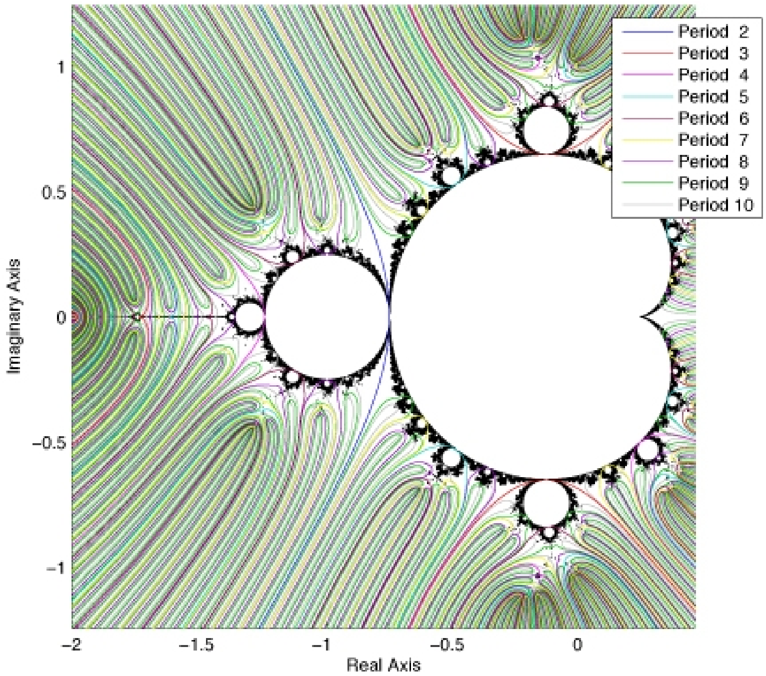

Figure 23: This image shows the Mandelbrot set along with all pairs of parameter rays which land at a root of a hyperbolic component of period at most 10. The parameter rays are coloured accordingly. In total, this gives 1966 rays. Each ray pair cuts the plane into two connected components. The component not containing zero, that is the one disjoint from the main cardioid, is the wake of the indifferent parameter at which the two rays land. This image was released into the public domain by its creator, Wikipedia user BeagleTheBagel.

Wakes are a useful construction because they allow us to portion off parts of M. It is natural, then, to define limbs.

Definition 4.21 Let c be a rationally indifferent parameter, and let $\mathcal{W}$ be the wake of c. Then the *limb with root* c is the set

$$\mathcal{W} \cap M$$

i.e. the part of M in the wake of c.

Fix H as any hyperbolic component of M. Recall that Theorem 3.15 gives us a parametrization γ_H of the boundary of H, which sends $e^{2\pi ip/q}$ to the unique parameter on the boundary of H with a rationally indifferent cycle whose multiplier is $e^{2\pi ip/q}$. Each such parameter has its own limb. Thus, there is a one-to-one correspondence between rational numbers in $\mathbb{Q}/\mathbb{Z}$, and limbs attached to H.

Specifically, for each hyperbolic component H of M, and for each $p/q \in \mathbb{Q}/\mathbb{Z}$, define $L(H, p/q)$ to be the unique limb of H whose root is the parameter $\gamma_H(e^{2\pi ip/q})$. We say $L(H, p/q)$ is the *p/q-limb of H*.

Example 4.22 Recall that H_0 denotes the unique hyperbolic component of M of period 1, and is the interior of the main cardioid. Looking back at Figure 22, we see that $L(H_0, 1/2)$ is the limb with root $-3/4$, since $-3/4 = \gamma_{H_0}(e^{2\pi i*(1/2)})$.

The motivation for linking rational numbers and limbs (given a fixed hyberbolic component) is that they provide an upper bound on the diameter of a limb.

Proposition 4.23 (The Yoccoz Inequality) *Let H be a hyperbolic component of M. Then there exists a constant k_H such that, for all $p/q \in \mathbb{Q}/\mathbb{Z}$ except $p/q = 0$, we have*

$$\text{diam } L(H, p/q) \leq \frac{k_H}{q}$$

Proof. See Proposition 4.2 of [Hubbard, 1993]. □

Keep in mind that for any $p/q \in \mathbb{Q}/\mathbb{Z}$, we are always taking p and q to be co-prime, else the above is nonsense. The crucial result of the proposition is that the size of the limbs gets arbitrarily small as q grows. We will use this to show local connectivity of M at all indifferent parameters, except for those where the corresponding multiplier is precisely 1. This corresponds to the case $p/q = 0$, the so-called primitive parabolic parameters.

We are, at last, able to establish MLC at more parameters. The below is an expansion of the proof in [Hubbard, 1993]. We detail the constructions of the proof as best we can, and will go through some examples and pictures of them too.

Theorem 4.24 *Suppose $f_c \in F$ has an indifferent cycle. Then M is locally connected at c.*

Proof. Let $(z_1, \ldots, z_m)$ be an indifferent cycle of f_c, of period m. Let ρ be the multiplier of the cycle. Then, since the cycle is indifferent, $\rho = e^{2\pi i t}$ for some $t \in \mathbb{R}/\mathbb{Z}$. Let $U \subseteq M$ be an arbitrary open set containing c, and let $\varepsilon > 0$ be such that $B(c, \varepsilon) \subseteq U$. We split into three cases.

Case 1. Suppose t is irrational. Then c belongs to the boundary of a unique hyperbolic component H, of period m. Let ρ_H be the conformal isomorphism $H \to B(0, 1)$, and $\rho_{\overline{H}}$ be the extension to the boundary, given by Theorem 3.15. Let k_H be the constant provided by Proposition 4.23.

For each $\delta > 0$, let $R(\delta) = \{r \in \mathbb{Q}/\mathbb{Z} : |t - r| < \delta\}$ be the set of rational numbers within δ of t. Let $Q(\delta) = \inf\{q : p/q \in R(\delta)\}$. Note that $Q(\delta) \to \infty$ as $\delta \to 0$, since better rational approximations of t will require larger denominators. Take $\delta_1 > 0$ such that $Q(\delta_1) > \frac{2k_H}{\varepsilon}$. Note that, by Proposition 4.23, we have diam $L(H, p/q) < \varepsilon/2$ for all $p/q \in R(\delta_1)$.

Next, for each $\delta > 0$, define $S(\delta)$ to be the segment of $D(0, 1)$ formed by drawing a line between $e^{2\pi i(t-\delta)}$ and $e^{2\pi i(t+\delta)}$. Certainly, diam $S(\delta) \to 0$ as $\delta \to 0$. We translate this segment from the unit ball to $\overline{H}$ using the map $\rho_{\overline{H}}$. Note that $c \in \rho_{\overline{H}}^{-1}(S(\delta))$ for all $\delta > 0$. Since $\rho_{\overline{H}}$ is a homeomorphism, we will have diam $\rho_{\overline{H}}^{-1}(S(\delta)) \to 0$ as $\delta \to 0$. Take $\delta_2 > 0$ such that diam $\rho_{\overline{H}}^{-1}(S(\delta_2)) < \varepsilon/2$.

Let $\delta_3 = \min\{\delta_1, \delta_2\}$. Form the set

$$N = \rho_{\overline{H}}^{-1}(S(\delta_3)) \cup \left(\bigcup_{p/q \in R(\delta_3)} L(H, p/q) \right)$$

The root point of every limb in N is less than $\varepsilon/2$ away from c, since diam $\rho_{\overline{H}}^{-1}(S(\delta_3)) < \varepsilon/2$. Also, every point in a given limb of N is less than $\varepsilon/2$ away from the limb's root point, since diam $L(H, p/q) < \varepsilon/2$ for all $p/q \in R(\delta_3)$. Therefore, every point in N is less than ε away from c, so $N \subseteq B(c, \varepsilon) \subseteq U$.

Finally, it may be shown that every point in a small enough neighbourhood of c is either in $\overline{H}$, or in a limb attached to H at some rationally indifferent parameter.[9] Since N consists of a) every limb within a certain distance of c, and b) a neighbourhood of c as a subspace of $\overline{H}$, we conclude that N is

[9]This is the so-called "no ghost limbs" theorem, see Section 4 of [Schleicher, 1999] for a proof and discussion of this result.

a neighbourhood of c in M. Finally it remains to show that N is connected. Certainly the limbs are all connected to the $\overline{H}$-neighbourhood of c, which itself is certainly connected. Finally, each individual limb itself must be a connected subset of $\mathbb{C}$, since otherwise M would not be connected. We conclude M is locally connected at c.

Case 2. Suppose $t = p/q \in \mathbb{Q}/\mathbb{Z}$ is non-zero. Then c belongs to the boundary of two hyperbolic components. One of them, say H_1, is of period m, and $c = \rho_{\overline{H_1}}(e^{2\pi it})$ (in the notation of **Case 1**). The other one, say H_2, is of period a multiple of m, and $c = \rho_{\overline{H_2}}(0)$ is the root of H_2. That is, c has internal angle t on the boundary of H_1, and H_2 is a smaller hyperbolic component than H_1, attached to H_1 at c.

Precisely as above, we may construct a set N_1 consisting of the 'segment' $\rho_{\overline{H_1}}^{-1}(S(\delta))$ of H_1 and all the attached limbs, such that the diameter of the segment is less than $\varepsilon/2$, and the diameters of all the limbs is less than $\varepsilon/2$. We do the same for H_2, constructing N_2 on the other side of c. Then $\mathrm{int}(N_1 \cup N_2)$ is a connected open set contained in U, which contains c, so that M is locally connected at c.

Case 3. The remaining case is $\rho = 1$, where c is a 'primitive' indifferent parameter. This is a special case and requires a deeper theory to tackle. Hubbard outlines a proof using the theory of *Mandelbrot-like families* in Theorem 14.6 of [Hubbard, 1993], and Tan Lei uses the techniques of *parabolic implosion* to provide a proof in the paper [Lei, 2000]. □

We demonstrate the construction of an arbitrary connected open neighbourhood of one irrationally indifferent parameter, and one (non-primitive) rationally indifferent parameter. The natural place to look is on ∂H_0, the boundary of the main cardioid, since this is algebraically easiest to work with.

In fact, we can obtain an explicit form for the paramaterization of ∂H_0. Let $c \in \partial H_0$. Then f_c has an indifferent fixed point z, so that $z^2 + c = z$, and $f'(z) = 2z = e^{2\pi it}$, where $t \in \mathbb{R}/\mathbb{Z}$ is the internal angle of c on ∂H_0. Eliminating z from this pair of equations yields

$$c = z - z^2 = \frac{1}{2}e^{2\pi it} - \frac{1}{4}e^{4\pi it} = \gamma_{H_0}(t)$$

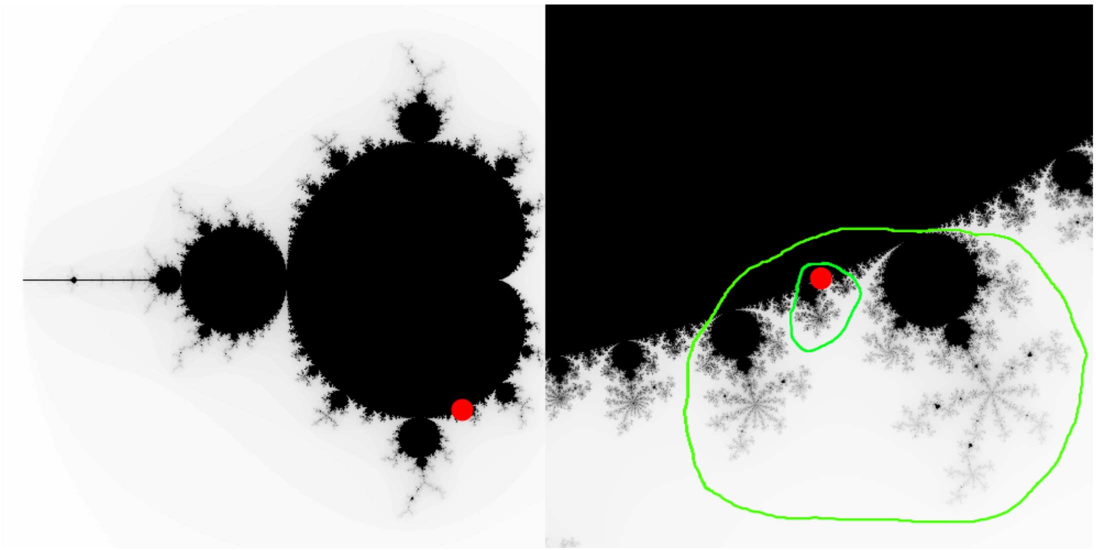

Figure 24: On the left, the Mandelbrot set is shown, with the approximate position of the irrationally indifferent parameter $c = \gamma_{H_0}(\sqrt{2}/2)$ indicated by the red dot. On the right, we zoom in on a neighbourhood of c. To construct a connected neighbourhood of c, we cut off a segment of the main cardioid (which is the unique hyperbolic component whose boundary c belongs to), and take this segment along with all of the limbs attached to the segment as the neighbourhood. Two such neighbourhoods are circled in green. Note that by taking a smaller segment in the smaller neighbourhood, we have a smaller largest limb. This construction generates arbitrarily small connected neighbourhoods of c, since there are limbs attached at every rational number between the internal angles of the outermost limbs of any given segment. Thus, smaller segments will have limbs whose internal angles are finer and finer rational approximations for $\sqrt{2}/2$, so by Proposition 4.23, such limbs get smaller and smaller.

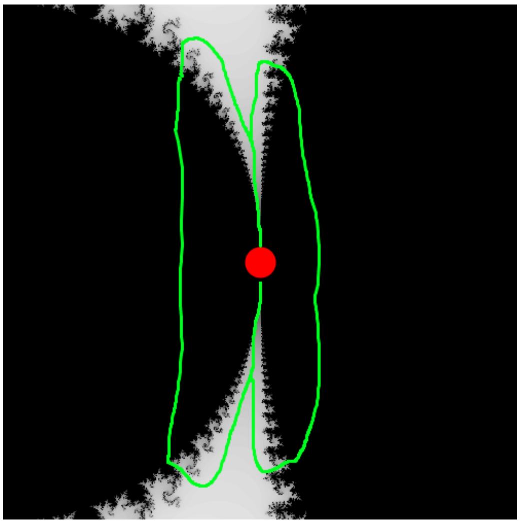

Figure 25: A neighbourhood of the rationally indifferent parameter $c = -3/4 = \gamma_{H_0}(-1)$. This parameter is the 'first period doubling bifurcation point' in M's period doubling cascade along the real axis. It is the unique parameter in the intersection of the boundaries of H_0 (the main cardioid) and H_1 (the unique hyperbolic component of period 2). To construct a connected neighbourhood of c, we cut off segments of both hyperbolic components, and take the union of both segments together with all the limbs that are attached to them both. The internal angles of the limbs of H_0 in smaller and smaller neighbourhoods of c are better and better 'rational approximations of -1', e.g. 0.9, 0.99, 0.999, ... The internal angles of the limbs of H_1 in smaller and smaller neighbourhoods of c are approximating 0, e.g. 0.1, 0.01, 0.001, ... By Proposition 4.23, the diameters of these limbs shrink as we take smaller segments, so that we can construct arbitrarily smaller connected neighbourhoods of c.

Yoccoz puzzles

Throughout this section, we need only concern ourselves with those $f_c \in F$ which have all repelling cycles. For, if there is an attracting cycle, then M is locally connected at c by Theorem 3.13, and if there is an indifferent cycle, then M is locally connected at c by Theorem 4.24. To this end, define $\widetilde{F}$ to be the subset of F consisting of those quadratics whose cycles are all repelling.

The approach of Yoccoz was to start in the dynamical plane, showing that when such a quadratic is not infinitely renormalizable, the Julia set is locally connected. We are actually able to explicitly construct arbitrarily small connected neighbourhoods at each point. The construction is known as the *Yoccoz puzzle.*

Now, since we are excluding hyperbolic and indifferent parameters from our considerations, we have no interest in the main cardioid H_0, which is disjoint from $\widetilde{F}$. As such, we look only at points that are in some limb $L(H_0, p/q)$ attached to the main cardioid at internal angle p/q. Suppose $c \in \widetilde{F} \cap L(H_0, p/q)$. Then f_c, as a quadratic, has two fixed points, which must both be repelling. One of the fixed points, traditionally denoted β, is the landing point of the ray $R_c(0)$ of external angle 0. Meanwhile, α is used to denote the other fixed point. α is characterised as the least repelling fixed point of f_c.[10] Since $c \in L(H_0, p/q)$, it must be the case that there are precisely q dynamical rays landing at α - this fact will be crucial in the construction of the Yoccoz puzzle. Label these rays $R_1, \ldots, R_q$.

The idea of the puzzle is to break the Julia set up into compact 'pieces' cut out by the rays, that join together like a jigsaw puzzle, and whose pre-images under f_c cover the Julia set by smaller and smaller connected sets. The 'puzzle pieces' are defined inductively, and so we must start with some open region that bounds the Julia set. Naturally, we look to the Green's function for this purpose. Take $r > 1$, and define

$$\mathcal{U}_0 = \{z \in \mathbb{C} : G_c(z) < r\}$$

so that $\mathcal{U}_0$ is the interior of the region bounded by the equipotential $\Gamma_c(\log r)$. By definition of the Green's function of $K(c)$, we have $K(c) \subset \mathcal{U}_0$. Note our choice of $r > 1$ makes no difference thus far, and indeed won't down the line either, since we will be shrinking $\mathcal{U}_0$ to arbitrarily small regions. Next, define

$$\mathcal{R}_0 = \bigcup_{i=1}^{q} (R_i \cap U_0)$$

as the set of all points which are in both one of the q rays landing at α, and within $\mathcal{U}_0$. We can now inductively define the puzzle of a quadratic.

Definition 4.25 Let $f_c \in \widetilde{F}$. The *puzzle* $\mathcal{P}_c$ of f_c is a pair of sequences

$$\mathcal{P}_c = (\mathcal{U}, \mathcal{R})$$

where $\mathcal{U} = (\mathcal{U}_0, \mathcal{U}_1, \ldots)$ and $\mathcal{R} = (\mathcal{R}_0, \mathcal{R}_1, \ldots)$ are defined inductively by $\mathcal{U}_n = f_c^{-1}(\mathcal{U}_{n-1})$ and $\mathcal{R}_n = f_c^{-1}(\mathcal{R}_{n-1})$.

That is, we look in on a part of the plane bounded by some equipotential, and take the parts of the q rays landing at α that are inside that part of the plane, and then the puzzle is formed as the sequences of pre-images of both the part of the plane, and the parts of the rays. Note that we necessarily have $\mathcal{U}_0 \supset \mathcal{U}_1 \supset \ldots$, and note also that at any point in the sequences, the rays in $\mathcal{R}_i$ cut $\mathcal{U}_i$ into open regions. This leads to our next definition.

Definition 4.26 Let $f_c \in \widetilde{F}$ have puzzle $\mathcal{P}_c$. For each $n \in \mathbb{N}$, the *puzzle pieces of* $\mathcal{P}_c$ *at depth* n are the elements of the set

$$\mathcal{P}_c(n) = \{\overline{X} : X \text{ is a connected component of } \mathcal{U}_n \backslash \mathcal{R}_n \text{ in } \mathcal{U}_n\}$$

[10] In fact, these observations hold for any $f_c \in F$, except the special case $c = 1/4$ where $\alpha = \beta$.

So, a puzzle piece is the closure of one of these q regions cut out from a connected component of some $\mathcal{U}_n$. Together, the puzzle pieces at a given depth n tile the set $\mathcal{U}_n$, i.e. they partition this set, with the only overlap being their boundaries. The regions $\mathcal{U}_n$ are cut out into pieces of a jigsaw puzzle by the pre-images of the dynamical rays landing at α. Since each $\mathcal{U}_n$ is contained in $\mathcal{U}_{n-1}$, we must have that every puzzle piece at depth n is contained in some puzzle piece at depth $n-1$.

The puzzle pieces are a sort of *Markov partition* for the dynamics of a quadratic - distinct puzzle pieces at the same depth have disjoint interiors, puzzle pieces of larger depth are contained in puzzle pieces of smaller depth, and f_c maps a puzzle piece at depth n to a puzzle piece at depth $n-1$.

By definition of the dynamical rays $R_1, \dots, R_q$ landing at α, the only point of intersection of the rays and $K(c)$ is α, so that $\mathcal{R}_0 \cap K(c) = \{\alpha\}$. Then since $K(c)$ is completely invariant under f_c (Proposition 2.10), the boundary of a puzzle piece will intersect $K(c)$ at some collection of pre-images of α. In fact, if $\alpha \notin O(c, 0)$, then each pre-image of α is in exactly q puzzle pieces at all sufficiently large depths (more precisely, on the boundary of q puzzle pieces).

We get an edge case out of the way now, when $\alpha \in O(c, 0)$, because we want to assume this previous observation will hold. Thankfully, this case was dealt with in [Douady and Hubbard, 1985a]. Recall that our current goal is to use puzzles to show local connectivity of Julia sets.

Proposition 4.27 *Let $f_c \in \widetilde{F}$, and suppose $\alpha \in O(c, 0)$. Then $J(c)$ is locally connected.*

Since we want to use the puzzle pieces as arbitrarily small connected neighbourhoods, we want to show that they get arbitrarily small. Accordingly, we introduce *ends*.

Definition 4.28 Let $f_c \in \widetilde{F}$ have puzzle $\mathcal{P}_c$. An *end* $\mathcal{E}$ of $\mathcal{P}_c$ is a nested sequence of puzzle pieces. That is, an end is any sequence of puzzle pieces $\mathcal{E} = (X_0, X_1, \dots)$ such that $X_0 \supset X_1 \supset \dots$

The set of all ends of $\mathcal{P}_c$ is denoted $\mathcal{E}_c$.

When we come to constructing a basis of connected neighbourhoods using puzzle pieces, we will essentially be dealing with some particular end. Thus we now want to find a way to show all ends get arbitrarily small.

Definition 4.29 The *impression* of an end $\mathcal{E} = (X_0, X_1, \dots)$ is the set

$$\mathcal{J}(\mathcal{E}) = \bigcap_{n=0}^{\infty} X_n$$

It may be shown that for any end $\mathcal{E}$, the impression $\mathcal{J}(\mathcal{E})$ is a compact, connected subset of $\mathbb{C}$. The impression can show us that an end gets arbitrarily small in the following way.

Definition 4.30 Let $\mathcal{E}$ be an end. Then we will say $\mathcal{E}$ *shrinks to a point* if the impression $\mathcal{J}(\mathcal{E})$ consists of a single point.

If an end $\mathcal{E} = (X_0, X_1, \dots)$ shrinks to a point, then the diameters of the sets X_n must necessarily tend to 0. Since it may also be shown that the intersection of any puzzle piece with the Julia set is connected, we can get a result on local connectivity.

Theorem 4.31 *Let $f_c \in \widetilde{F}$, and suppose that every end $\mathcal{E} \in \mathcal{E}_c$ shrinks to a point. Then $J(c)$ is locally connected.*

Proof. If the fixed point $\alpha \in O(c, 0)$, we get the result from Proposition 4.27. Otherwise, all pre-images of α belong to q puzzle pieces at sufficiently large depths. Let $z \in J(c)$. Then there necessarily exists an end $\mathcal{E} = (X_0, X_1, \dots)$ such that $z \in X_n$ for all $n \in \mathbb{N}$. Now, either z is a pre-image of α, or it isn't. If it is not, then z cannot be on the boundary of any of the puzzle pieces X_n, so instead is in the interior of each. In this case, the sets $X_n \cap J(c)$ are connected neighbourhoods of z whose diameters approach 0, so that there are arbitrarily small connected neighbourhoods of z in $J(c)$, and $J(c)$ is locally connected at z. Finally, if z is a pre-image of α, then at sufficiently large depths, z is on the boundary of q puzzle pieces. Let Y_n be the union of the q puzzle pieces at depth n which contain z, where n is sufficiently large. Each puzzle piece in Y_n is in some end, which shrinks to a point, so that the diameters of each puzzle piece in Y_n tends to 0 as $n \to \infty$. Therefore, the diameter of Y_n tends to 0 also. Thus, just as before, the sets $Y_n \cap J(c)$ are arbitrarily small connected neighbourhoods of z in $J(c)$, and $J(c)$ is again locally connected at z. □

It would remain to show that all ends of non-infinitely renormalizable quadratics (with all cycles repelling) must shrink to points. This involves diverging sums of *moduli of annuli* (see Section 4.4) of nested puzzle pieces, and *tableaux of ends*, and other technical machinery. After that, it remains still to transfer the results to the parameter plane. This is done with an analogous construction called the *para-puzzle*, which provides a basis of connected neighbourhoods in M for all parameters c whose quadratic is not infinitely renormalizable. The arbitrarily small shrinking of the para-puzzles is demonstrated using the similarity of $J(c)$ and M near c. All of these details may be found in [Hubbard, 1993]. After all this, we are able to deduce Theorem 4.17.

Example 4.32 As a demonstration of the Yoccoz puzzle, we look at the puzzle $\mathcal{P}_\omega$ for the Douady Rabbit quadratic f_ω. We choose our initial region $\mathcal{U}_0$ for the puzzle to be the interior of the region bounded by $\Gamma_\omega(1.2)$. The α fixed point lies at $z = \frac{1-\sqrt{1-4\omega}}{2} \approx -0.276 + 0.48i$, and has 3 dynamical rays landing at it. The puzzle pieces at depth 0 can then be formed.

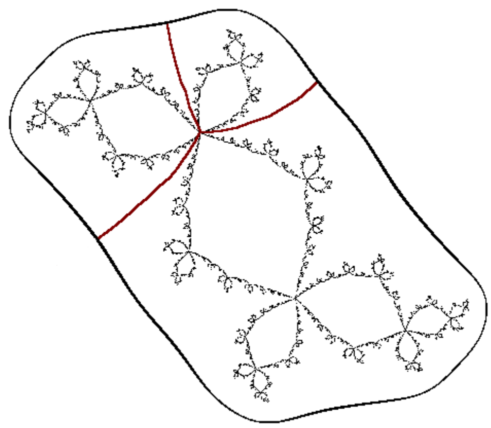

Figure 26: The puzzle pieces at depth 0 of $\mathcal{P}_\omega$ are the 3 closed regions cut out by the 3 dynamical rays landing at α, and bounded by the equipotential $\Gamma_\omega(1.2)$.

The pre-image of $\mathcal{U}_0$ is the interior of the region bounded by the equipotential $\Gamma_\omega(\sqrt{1.2})$. This is our $\mathcal{U}_1$. There are 6 'pre-images rays', 3 of which are the rays landing at α, and the other 3 land at $-\alpha$. The puzzle pieces at depth 1 can then be formed.

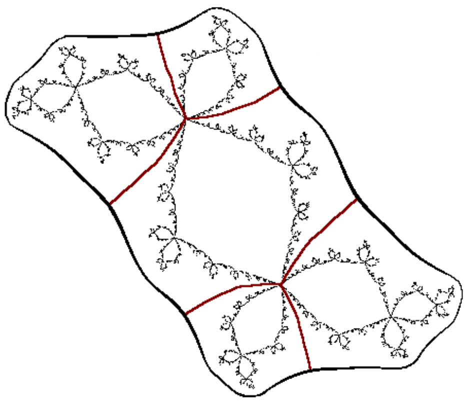

Figure 27: The puzzle pieces at depth 1 of $\mathcal{P}_\omega$ are the 5 closed regions cut out by the 6 pre-images of the 3 dynamical rays landing at α, and bounded by the equipotential $\Gamma_\omega(\sqrt{1.2}) = f_\omega^{-1}(\Gamma_\omega(1.2))$.

As an alternative example to see how the same ideas can be applied to any Julia set, we present the following illustration which is Figure 8.3 on page 133 of [McMullen, 1994a].

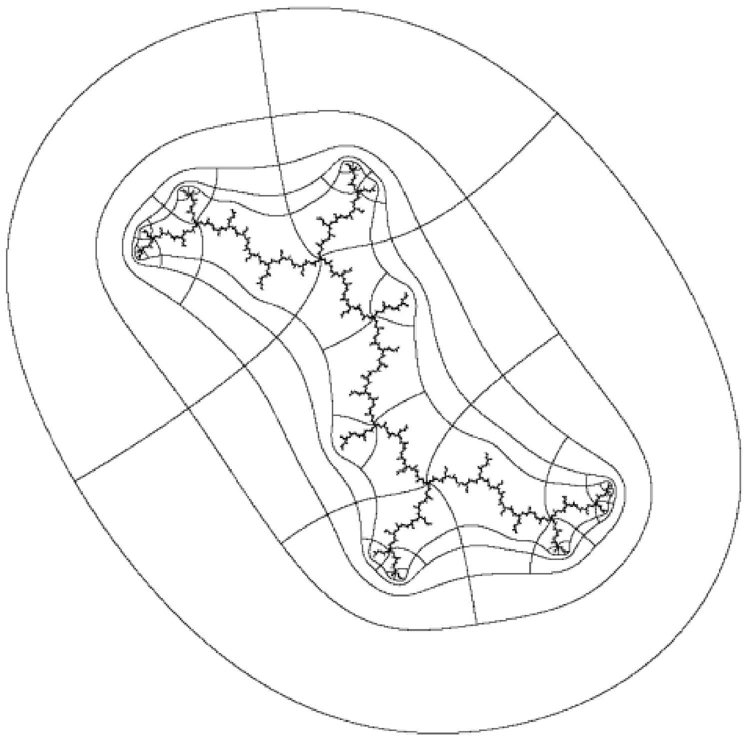

Figure 28: The puzzle pieces at depth 0, 1, 2, 3 and 4 of $\mathcal{P}_i$ are all shown together. That is, the puzzle pieces for $f_i : z \mapsto z^2 + i$.

In both of these examples, we may continue to go to deeper depths of the puzzle. The puzzle pieces at each depth will always tile the filled-in Julia set, and in the cases where all cycles are repelling and the map is not infinitely renormalizable, all ends of the puzzle will shrink to points, so that the puzzle pieces will get arbitrarily small at arbitrarily large depths. We will have each point in $J(c)$ either being a pre-image of the α fixed point and thus on the shared boundary of q pieces (where q is the number of rays landing at α), or otherwise being in the interior of puzzle pieces at all depths. We can see the phenomena clearly in the above 2 examples (which both have $q = 3$).

4.4 The *a priori bounds*, and quadratics that are infinitely renormalizable of bounded type

We are going to look briefly at one other published result that made progress towards the full MLC conjecture. Specifically, we look at [Kahn, 2006] and define the class of parameters for which local connectivity was shown in that work. We also define a property that infinitely renormalizable maps can possess, the so-called *a priori bounds*, which implies local connectivity of M at the corresponding parameter. It is through a priori bounds that local connectivity of M has often been demonstrated, with almost every entry in the above timeline utilising them. Notably the most recent result [Dudko and Lyubich, 2019] is somewhat of an exception - in that work, the class of parameters for which MLC is proven consists of some that have a priori bounds, and some that don't.

First we must discuss *annuli* in the complex plane.

Definition 4.33 Let $a \in \mathbb{C}, 0 < r < R$. Then, the *annulus with centre a, inner radius r, and outer radius R* is the set

$$A(a, r, R) = \{z \in \mathbb{C} : r < |z - a| < R\}$$

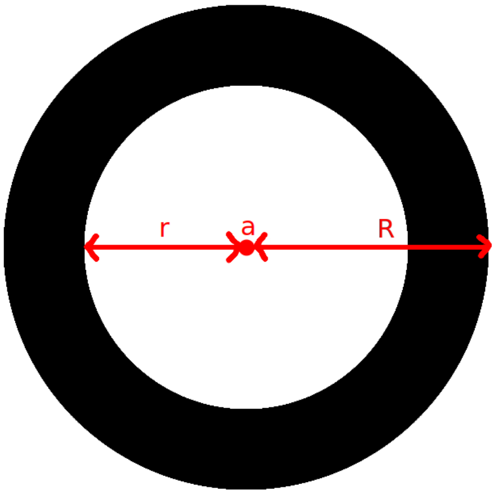

Figure 29: The annulus $A(a, r, R)$.

It is easy to conformally map an annulus to an annulus of the standardised form $A(0, 1, R)$ for some $R > 1$. Annuli are very regular subsets of the complex plane, and for our applications, it is far more typical that we will be dealing with 'annuli-like' subsets of the plane. We have already defined the machinery required for this: *conformal equivalence*. In fact, we need only that a given subset is homeomorphic to an annulus to get conformal equivalence.

Proposition 4.34 *Let U be an open subset of $\mathbb{C}$, and suppose U is homeomorphic to some annulus. Then U is conformally equivalent to $A(0, 1, R)$ for some unique $R > 1$.*

The uniqueness allows for the following definition, which provides a notion of size for these annuli-like subsets.

Definition 4.35 Let U be an open subset of $\mathbb{C}$ that is conformally equivalent to $A(0, 1, R)$ for some $R > 0$. Then the *modulus of U* is

$$\text{mod}(U) = \frac{\log R}{2\pi}$$

There is a wide variety of places in complex dynamics, and in fact generally in complex analysis, in which we meet subsets that are conformally equivalent to annuli. In fact, anywhere where we meet nested subsets that are each conformally equivalent to $B(0,1)$, which as discussed previously, just means nested subsets of the complex plane that are simply connected. Then, the set-theoretic difference of successive nested such subsets will be conformally equivalent to annuli.

One such situation is the nested puzzle pieces discussed in Section 4.3, and indeed the shrinking of puzzle pieces was demonstrated by Yoccoz using the modulus of the sets formed by the set-theoretic difference of successive nested puzzle pieces. Though we don't look at this in any detail, it does serve as a motivating example, as we will discuss in a moment.

There is another situation in which we have already met nested subsets that are conformally equivalent to $B(0,1)$:

Proposition 4.36 *Let U, V be open subsets of $\mathbb{C}$ that are conformally equivalent to $B(0,1)$, and which satisfy $\overline{U} \subset V$. Then $V \backslash U$ is homeomorphic to an annulus, and thus conformally equivalent to $A(0,1,R)$ for some $R > 0$. In particular, if $f : U \to V$ is a quadratic-like mapping, then $V \backslash U$ is conformally equivalent to $A(0,1,R)$ for some $R > 0$.*

Suppose $f_c \in F$ is infinitely renormalizable. Then there is a monotone strictly increasing sequence $(p_0, p_1, p_2, \ldots)$ of integers such that f_c is renormalizable at level p_k for each $k \in \mathbb{N}$, and thus a sequence of quadratic-like mappings

$$f_c^{\circ(p_k)} : U_k \to V_k$$

In fact, there are various such sequences of integers. We will call any such sequence an *infinite sequence of renormalization periods*. We are now able to define a priori bounds.

Definition 4.37 Let $f_c \in F$ be infinitely renormalizable. Then f_c is said to *have/enjoy a priori bounds* if there exists an infinite sequence of renormalization periods $(p_0, p_1, p_2, \ldots)$, and an associated sequence of quadratic like mappings

$$f_c^{\circ(p_k)} : U_k \to V_k$$

such that

$$\text{mod}(V_k \backslash U_k) \geq \varepsilon$$

for each $k \in \mathbb{N}$, and for some universal lower bound $\varepsilon > 0$.

In [Lyubich, 1997], Mikhail Lyubich describes a priori bounds as "a basic geometric quality of infinitely renormalizable maps" and "a key to the renormalization theory, problems of rigidity and local connectivity". They were first introduced by Dennis Sullivan in the work [Sullivan, 1988].

As we mentioned above, a motivating example for the definition of a priori bounds is nested sequences of puzzle pieces, which give rise to an infinite sequence of smaller and smaller annuli-like subsets. In particular, it is crucial to the proof of Theorem 4.17 to show that the infinite sum of the moduli of each annuli-like subset is infinite, and this is something of a theme in results relating to the MLC conjecture. It is clear that the infinite sum of moduli of the annuli-like subsets formed in Definition 4.37 will be infinite when there is a priori bounds. This has the following far-reaching consequence, which gives a priori bounds their importance.

Theorem 4.38 *Let $f_c \in F$ be infinitely renormalizable, and suppose f_c has a priori bounds. Then $J(c)$ is locally connected, and M is locally connected at c.*

Our final act of this section is to define the quadratics that are *infinitely primitively renormalizable of bounded type* - for whom a priori bounds was proven by Jeremy Kahn in [Kahn, 2006]. This requires two new definitions.

Definition 4.39 Suppose $f_c \in F$ is renormalizable, with $f = f_c^{\circ n}|_U : U \to V$ being the quadratic-like mapping used in some renormalization. Let $K(f)$ be the quadratic-like map's Julia set. Then, the renormalization is *primitive* if the elements of $\{f_c^{\circ k}(K(f)) : k = 0, 1, \ldots, n-1\}$ are pairwise disjoint. An infinitely renormalizable map is *infinitely primitively renormalizable* if it has infinitely many primitive renormalizations.

Example 4.40 In Example 4.13, we looked at the renormalizable map f_c, with $c \approx -1.7577 + 0.0134i$. In fact, the renormalization in that instance was primitive, since as discussed in that example, the orbit jumped between 3 'mini-Rabbits', as shown in Figure 20. The collection of these 3 mini-rabbits is precisely the collection whose pairwise empty intersection makes the renormalization a primitive one.

Example 4.41 In Example 4.16, we looked at the renormalizable *Feigenbaum map* f_{c_F}, with $c_F = -1.4011551890$. In this case, letting f be the quadratic-like map used in the renormalization, we have the collection $\{K(f), f_{c_F}(K(f))\}$ that will tell us whether or not the renormalization is primitive. In fact, these two sets are not disjoint, and they touch at a single common point touch at a single common point, so that the renormalization is not primitive. Whenever primitiveness fails for some $f_c^{\circ n} : U \to V$ because of this non-empty intersection, the intersection will always consist of a single point, which will be a repelling fixed point of $f_c^{\circ n}$.

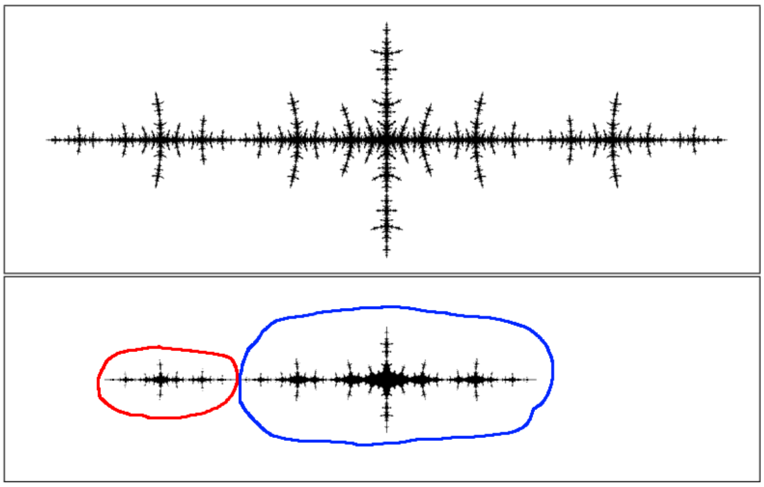

Figure 30: The top image shows the Julia set $J(c_F)$. The bottom image shows $K(f)$ circled in blue, and $f_{c_F}(K(f))$ circled in red. The pair touch at a single point, which is a repelling fixed point of $f_{c_F}^{\circ 2}$. This image is taken from [McMullen, 1994a], where it is Figure 7.5 on page 114.

Definition 4.42 Let $f_c \in F$. Then f_c is *infinitely primitively renormalizable of bounded type* if there exists an infinite sequence of renormalization periods $(p_0, p_1, p_2, \ldots)$ such that f_c is primitively renormalizable at level p_k for each $k \in \mathbb{N}$, and

$$p_{k+1}/p_k \leq B$$

for each $k \in \mathbb{N}$, and for some universal upper bound $B > 0$.

At last, we are able to state the result of [Kahn, 2006].

Theorem 4.43 *Suppose $f_c \in F$ is infinitely primitively renormalizable of bounded type. Then f_c has a priori bounds.*

Corollary 4.44 *Suppose $f_c \in F$ is infinitely primitively renormalizable of bounded type. Then $J(c)$ is locally connected, and M is locally connected at c.*

5 The Pinched Disc Model of the Mandelbrot Set

An author of a project on the MLC conjecture is spoiled for choice when it comes to writing about the consequences of the conjecture. The difficulty comes in finding a way to explain any such consequence in a vaguely intelligible way, given that we hope for a highly-attuned undergraduate student being able to understand what we present. If MLC were to be proven, a whole host of other conjectures about the dynamics of the quadratics would immediately follow, for example various rigidity conjectures and the *no invariant line fields* conjecture for quadratic Julia sets. Sadly we won't be covering any of this. The survey [Benini, 2018] has excellent exposition on all of the aforementioned.

Instead, our final chapter explores one particular way that local connectivity of the Mandelbrot set could be applied. The landing of all parameter rays, and the associated parametrisation of the boundary (both given by the Carathéodory-Torhorst Theorem 3.28) allows us to build a topological space called the *pinched disc model* of M. In fact, we are able to construct the pinched disc model even without local connectivity, either using the landing of all parameter rays of rational external angle, or inducing the model from similar models for Julia sets. In any case, the model is homeomorphic to M if and only if M is locally connected.

If the constructed space is indeed homeomorphic to M, then we may use it to prove results about M and therefore about the dynamics of the quadratics. We will even sketch a proof of how this idea can be used to show that the MLC conjecture implies density of hyperbolicity for complex quadratic maps.

We present here the construction of the pinched disc model via the *Quadratic Minor Lamination.* We rely heavily on the sources [Thurston, 1985] (and its appendix [Schleicher, 2009]), and [Douady, 1993].

5.1 Geodesics, convex hulls, and hyperbolic geometry

Since the definition of a *lamination* requires it, we begin by familiarising ourselves with notions of Euclidean and hyperbolic geometry on the closed unit disc.[11]

Although the definitions that are given below extend naturally to any number of contexts, we give them in the rather narrow context of the unit disc $D(0,1)$, since this is all that an explanation of the pinched disc model demands. We begin in the familiar Euclidean setting.

Definition 5.1 A *Euclidean geodesic* in $D(0,1)$ is a straight line between any two points $x, y \in D(0,1)$. We denote the Euclidean geodesic between x and y by $\mathrm{EG}(x,y)$.

Clearly, for $x, y \in D(0,1)$, the Euclidean geodesic $\mathrm{EG}(x,y)$ gives the shortest path between x and y. Indeed, it is not difficult to explicitly determine this geodesic as $\mathrm{EG}(x,y) = \{(1-t)x + ty : t \in [0,1]\}$.

Definition 5.2 Let $E \subseteq D(0,1)$. Then E is *Euclidean convex* if $\mathrm{EG}(x,y) \subseteq E$ for all $x, y \in E$. The *Euclidean convex hull* of E is the smallest convex set $C \subset D(0,1)$ such that $E \subseteq C$. We denote the Euclidean convex hull of E by $\mathrm{ECH}(E)$.

Note that $\mathrm{ECH}(E)$ is well-defined for all $E \subseteq D(0,1)$, since $D(0,1)$ is itself Euclidean convex. We can think of convex sets as being those in which we are always able to take the most direct route between any two points, while staying in the set.

Example 5.3 Let $x, y \in D(0,1)$, and let $g = \mathrm{EG}(x,y)$. Then $\mathrm{ECH}(g) = g$. More generally, any point, line or polygon in $D(0,1)$ is its own convex hull.

[11] We note now that the adjective hyperbolic is getting a bit overloaded, given earlier and future discussions of hyperbolic dynamics, hyperbolic maps, density of hyperbolicity etc. It is important to keep in mind that this 'hyperbolic' is distinct from the one referenced in *hyperbolic geometry.*

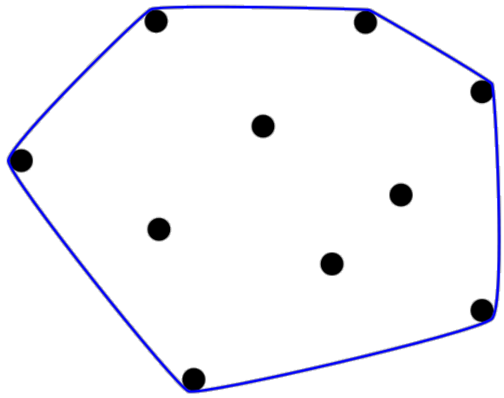

Figure 31: A collection of points is shown. Let X be the set of all these points. Then the convex hull $\mathrm{ECH}(X)$ of X is the region bounded by the blue polygon. This image was created by Wikipedia user Maksim, and released into the public domain by them.

The above is, technically, all that we need - we could stop here, and define laminations, the quadratic minor lamination, and the pinched disc model using Euclidean geodesics and Euclidean convex hulls, omitting any part of hyperbolic geometry. However, the theory goes through just the same if we instead use the hyperbolic geometry equivalents of the above definitions, and this makes for far more comprehensible pictures, as well as bringing out a nice visual aspect of the relationship between laminations and external rays. Thus our next step is to define hyperbolic geodesics, and hyperbolic convex hulls. We hope that it was useful to first define the Euclidean versions, as clear and very familiar analogues to the hyperbolic versions that we will use.

For thousands of years, the geometry encountered in mathematics was that defined by Euclid in his *Elements*. Euclid started from 10 fundamental axioms and postulates of geometry, and used these to prove hundreds of propositions and theorems. Euclid's methods were immensely successful, and his axiomatic approach set the standard for the millenia of mathematics that followed. The 5th and final postulate, called the *parallel postulate*, stated: given a line in the plane and a point not on the line, there is exactly one line through the point which does not intersect the first line[12]. The parallel postulate become something of a controversial matter - it stood out among the axioms and postulates as the most complicated (contrast with Axiom 5: *"The whole is greater than the part"*), and mathematicians generally believed that it should be possible to prove the parallel postulate from the other 9 axioms and postulates. A proof, however, never materialised, and the list of mathematicians who attempted a proof and failed ranges from Proclus of Constantinople (c. 410 A.D.) to Adrien-Marie Legendre (c. 1800 A.D.).

Eventually, it was shown that the parallel postulate does not necessarily follow from the other axioms and postulates. Furthermore, we may even get rid of the parallel postulate, and replace it with something else, and in the process obtain a whole new geometry, which is logically consistent, and perhaps even rather elegant. *Hyperbolic geometry* is what we obtain when we instead use the following axiom: given a line in the plane and a point not on the line, there are at least two lines through the point which do not intersect the first line.

We will work with hyperbolic geometry only in the unit disc, and this allows us to think in terms of *orthogonal circles*.

Definition 5.4 Let C, D be circles in the plane which intersect in two points. Then C and D are orthogonal if, at both points of intersection, the tangents of the circles are at a right angle to one another.

[12] Actually, this is *Playfair's axiom*, the modern (relative to Euclid, at least!) equivalent of the parallel postulate.

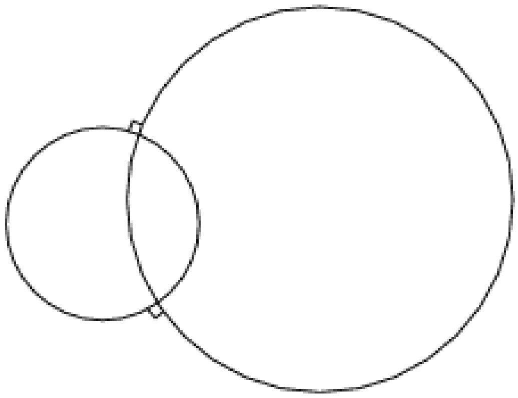

Figure 32: Two orthogonal circles. Supposing for a moment that the larger circle is $C(0,1)$, then the part of the smaller circle that is contained in $D(0,1)$ is precisely the hyperbolic geodesic between the two points of intersection of the circles. This image was taken from the Wolfram MathWorld entry on orthogonal circles, written by Eric W. Weisstein:
https://mathworld.wolfram.com/OrthogonalCircles.html

Definition 5.5 A *hyperbolic geodesic* in $D(0,1)$ is a curve between two points $x, y \in D(0,1)$ such that the curve is a subset of some circle that is orthogonal to $C(0,1)$, and that the curve is contained in $D(0,1)$. We denote the hyperbolic geodesic between x and y by $\mathrm{HG}(x,y)$.

While this definition more than suffices for our purposes, there are a couple of unresolved issues with it. First of all, it does not seem immediately clear that there is a unique hyperbolic geodesic between any $x, y \in D(0,1)$. Secondly, it makes the relationship between Euclidean and hyperbolic geodesics seem obscure. We briefly attempt to resolve these. On $D(0,1)$, it is possible to define a 'hyperbolic distance function', which gives rise to a hyperbolic geometry on $D(0,1)$. For this distance function, there is a unique shortest path between any two points $x, y \in D(0,1)$, and this path is precisely the hyperbolic geodesic as defined above. Thus, $\mathrm{EG}(x,y)$ is well-defined for any $x, y \in D(0,1)$, and the hyperbolic geodesic between x and y gives the shortest path between x and y in hyperbolic space, just as the Euclidean geodesic gives the shortest path between x and y in Euclidean space.

We will be exclusively concerned with hyperbolic geodesics between points on the boundary of $D(0,1)$. The following is a fun tool I've made on Desmos for getting an intuition for these objects!

https://www.desmos.com/calculator/robo02fsnr

As in the Euclidean case, we are able to define convexity.

Definition 5.6 Let $E \subseteq D(0,1)$. Then E is *hyperbolic convex* if $\mathrm{HG}(x,y) \subseteq E$ for all $x, y \in E$. The *hyperbolic convex hull* of E is the smallest hyperbolic convex set $C \subseteq D(0,1)$ such that $E \subseteq C$. We denote the hyperbolic convex hull of E by $\mathrm{HCH}(E)$.

Just as before, $\mathrm{HCH}(E)$ is well-defined for all $E \subseteq D(0,1)$, since $D(0,1)$ is itself hyperbolic convex.

Before moving on to laminations, we reiterate that we could freely use either Euclidean or hyperbolic geodesics/convex hulls where we need them, and that it is advantageous to use the hyperbolic ones. For this reason, we will drop the adjective 'hyperbolic' for these objects, i.e. simply say 'geodesic' for hyperbolic geodesics in $D(0,1)$, or 'convex hull' for the hyperbolic convex hull of a set.

5.2 Laminations, and the pinched disc model for Julia sets

Consider the geodesic $l = \mathrm{HG}(x,y)$ for $x, y \in D(0,1)$. It is defined in a purely geometric sense, but we are more interested in it as a set-theoretic object. That is, l is a set consisting of all the points on the shortest path between x and y. This will be important to keep in mind.

Definition 5.7 A *proto-lamination* is a set L of geodesics in $D(0,1)$, such that for all $l \in L$, the endpoints of l belong to $C(0,1)$.

Thus, a proto-lamination consists of geodesics that are 'hyperbolic chords' of $C(0,1)$ in the traditional sense of chords of a circle. The geodesics in a proto-lamination are called *leaves*. Since the geodesics themselves are sets, a proto-lamination is a set of sets. Recall the 'big cup' notation $\bigcup \mathbb{S}$ for a set of sets $\mathbb{S}$, which denotes the set of all elements which are in some element of $\mathbb{S}$.

Definition 5.8 A *lamination* is a proto-lamination L such that:

(i) Distinct leaves of L can intersect only at their endpoints, and

(ii) $C(0,1) \cup (\bigcup L)$ is closed in $D(0,1)$.

Figures 33, 34, 35 and 36 (illustrations by Ella Matza) are various examples of collections of geodesics in $D(0,1)$ which do or do not form laminations.

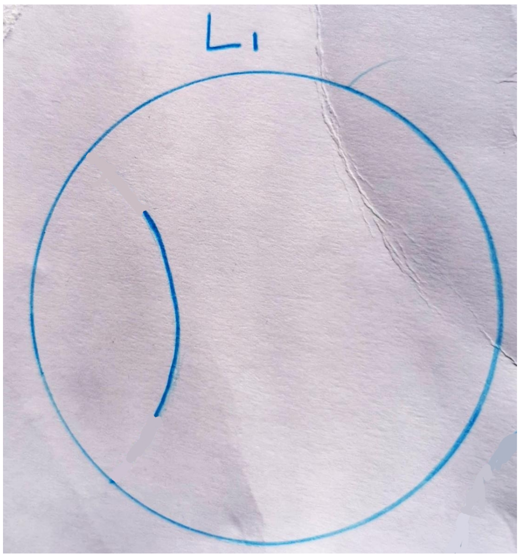

Figure 33: The set L_1 consisting of the above single geodesic does not form a proto-lamination and therefore not a lamination, since the endpoints of the geodesic are in $B(0,1)$ instead of the boundary $C(0,1)$.

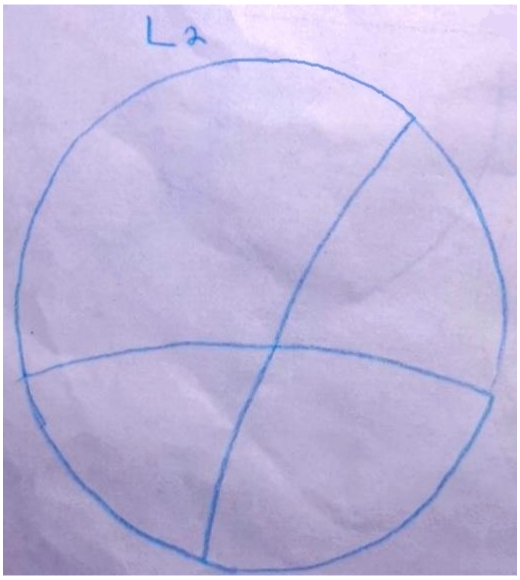

Figure 34: The collection L_2 of these two leaves forms a proto-lamination, but it is not a lamination since the two leaves intersect in $B(0,1)$.

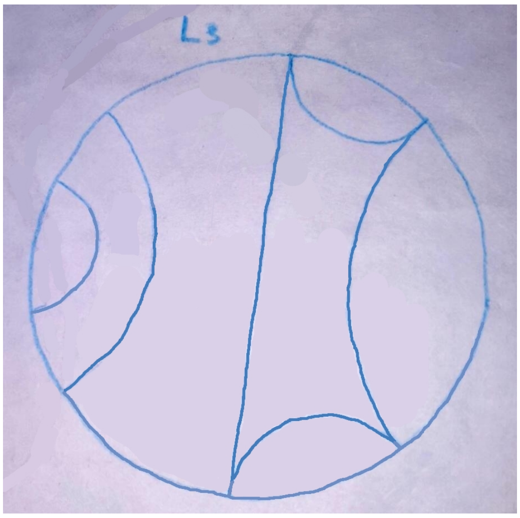

Figure 35: The collection L_3 of these 6 leaves forms a lamination. The endpoints of each leaf are on $C(0,1)$, and the leaves intersect each other only on $C(0,1)$ at their endpoints. It is closed since there are only finitely many leaves.

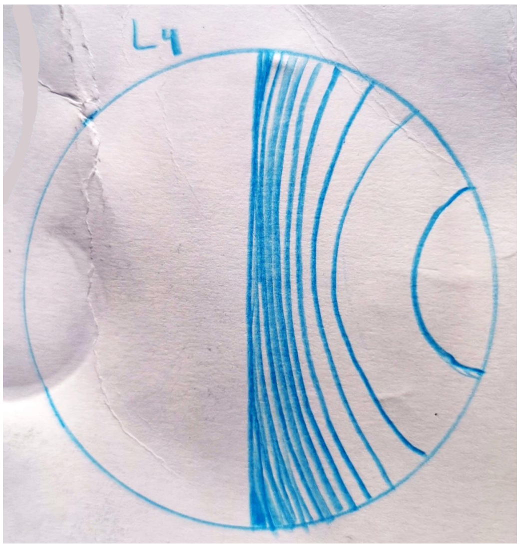

Figure 36: Let L_4 be the proto-lamination defined as follows. Choose a strictly monotonically increasing sequence $(x_n)_{n\in\mathbb{N}}$ such that $\{x_n : n \in \mathbb{N}\} \subset (0, 1/4)$ and $x_n \to 1/4$ as $n \to \infty$. For each $n \in \mathbb{N}$, let $a_n = e^{2\pi i x_n}$, $b_n = e^{2\pi i(1-x_n)}$, and define the leaf $l_n = \mathrm{HG}(a_n, b_n)$. L_4 is illustrated above. The leaves are pairwise disjoint, but L_4 does not contain the 'limit leaf' $l = \mathrm{HG}(i, -i)$, which is the vertical diameter of $D(0,1)$, and thus $C(0,1) \cup (\bigcup L_4)$ is not closed in $D(0,1)$. The set $L_4 \cup \{l\}$ is a lamination.

Let L be a lamination. A *gap* of L is the closure of a connected component of $D(0,1) \setminus (\bigcup L)$. That is, we take the unit disc and slice it up by cutting through all the leaves with our mathematical scissors. The gaps of L are the closures of the resulting cut-up pieces. Let $\mathcal{G}(L)$ denote the set of all gaps of L. For a gap $G \in \mathcal{G}(L)$, we may define the *periphery* of G as $P(G) = G \cap C(0,1)$, the set of points where the gap touches the boundary of $D(0,1)$.

For the rest of this section, reserve f to be the function $f : C(0,1) \to C(0,1)$ given by $f(z) = z^2$. It is through this map that we find something of a connection between particular laminations and quadratic dynamics. The particular laminations satisfy certain properties regarding the map f, as the next definition outlines.

Definition 5.9 A lamination L is *quadratic invariant* if the following all hold:

(i) (Forward invariance) For all $\mathrm{HG}(x,y) \in L$, we have either $\mathrm{HG}(f(x), f(y)) \in L$, or $f(x) = f(y)$

(ii) (Backward invariance) For all $\mathrm{HG}(x,y) \in L$, we have that $\mathrm{HG}(x_1, y_1)$ and $\mathrm{HG}(x_2, y_2)$ are distinct leaves of L, where $f^{-1}(x) = \{x_1, x_2\}$ and $f^{-1}(y) = \{y_1, y_2\}$

(iii) (Gap invariance) For all gaps $G \in \mathcal{G}(L)$, the image under f of the periphery of G has a convex hull which is either a gap, a leaf or a single point.

So, forward invariance means that the endpoints of a leaf either map to the endpoints of another leaf or collapse under f to a point. Backward invariance means that the pre-images of the endpoints of every leaf give us two of the leaves of L. Gap invariance means that, we find a way to map a gap by mapping its periphery, and then we come out of the periphery using the convex hull, and that for any gap, this process gives a gap, a leaf, or a point.

Let L be a quadratic invariant lamination, and let $\mathcal{L}$ be the the union of all leaves and gaps of L. Then the map f may be extended to a map

$$\widetilde{f} \ : \ C(0,1) \cup \mathcal{L} \ \rightarrow \ C(0,1) \cup \mathcal{L}$$

We start with the leaves, using forward invariance: let $\mathrm{HG}(x,y) \in L$. If $f(x) = f(y)$, then set $\widetilde{f}(t) = f(x)$ for all $t \in \mathrm{HG}(x,y)$. Otherwise, we simply map each $\mathrm{HG}(x,y)$ in a linear fashion to $\mathrm{HG}(f(x), f(y))$. The gaps are trickier. Let $G \in \mathcal{G}(L)$. By gap invariance, we obtain either a gap $H \in \mathcal{G}(L)$, a leaf $\mathrm{HG}(x,y) \in L$, or a point $z \in C(0,1)$. The single point case is easy, as we can map every point in the gap to z. We do not delve too deeply into the other cases, except for saying that the image of the gap as a whole will be either H or $\mathrm{HG}(x,y)$, depending on which one exists!

Suppose we have a quadratic invariant lamination L which has some leaf $\mathrm{HG}(x,y)$. Since $x, y \in C(0,1)$, we may write

$$x = e^{2\pi i s}, \qquad y = e^{2\pi i t}$$

for $s, t \in \mathbb{R}/\mathbb{Z}$. Recall that the distance between s and t in $\mathbb{R}/\mathbb{Z}$ is given by $\min\{|t-s|, |1-t-s|\}$, and is the smallest distance you would need to traverse from x to y via $C(0,1)$. We define the *boundary length* of $\mathrm{HG}(x,y)$ to be this distance. Note that the boundary length is always less than or equal to 1/2.

Proposition 5.10 *Let L be a quadratic invariant lamination, and define M to be the supremum of the boundary length, taken over all of L's leaves. Then there exists $l \in L$ that has boundary length M. That is, L has a maximal leaf with respect to boundary length.*

Any leaf $l \in L$ which has maximal boundary length is called a *major* leaf. If $M = 1/2$, then l is a diameter of $C(0,1)$, and thus is the unique major leaf of L. If $M \neq 1/2$, then L has precisely 2 major leaves.

Proposition 5.11 *Let L be a quadratic invariant lamination with two major leaves m_1 and m_2. Then $\widetilde{f}(m_1) = \widetilde{f}(m_2)$.*

The unique leaf in the image of the major leaves is called the *minor* leaf. It's possible for the minor leaf to be a point and not a leaf, in which case we say it is a *degenerate* minor leaf.

The reason for our interest in quadratic invariant laminations is that they can be used to build topological models, not only of the Mandelbrot set, but of Julia sets too. Suppose $f_c \in F$ with $J(c)$ locally connected. Then by Corollary 3.29, we have a parametrisation $\gamma_c : C(0,1) \rightarrow J(c)$ of $J(c)$, which is continuous and surjective, and for each $t \in \mathbb{R}/\mathbb{Z}$, the dynamical ray $R_c(t)$ lands at the point $\gamma_c(e^{2\pi i t})$.

We will use the landing of all the dynamical rays to associate f_c, or more accurately $K(c)$, with a lamination Λ_c. For each $z \in J(c)$, there is a finite number of rays landing at z, say $R_c(t_1), \ldots, R_c(t_k)$. Let $z_1 = e^{2\pi i t_1}, \ldots, z_k = e^{2\pi i t_k}$, and define

$$\lambda_c(z) = \{\mathrm{HG}(z_1, z_2), \mathrm{HG}(z_2, z_3), \ldots, \mathrm{HG}(z_{k-1}, z_k), \mathrm{HG}(z_k, z_1)\}$$

Then we may define the lamination

$$\Lambda_c = \bigcup_{z \in J(c)} \lambda_c(z)$$

So, if there are two dynamical rays landing at a point of the Julia set, the lamination has a leaf whose endpoints on the circle are given by the external angles of the rays. If there are many dynamical rays landing at a point, we do the same thing, but with each consecutive pair of rays.

Proposition 5.12 *Let $c \in M$ be such that $J(c)$ is locally connected. Then the above construction Λ_c is a quadratic invariant lamination.*

The minor leaf of Λ_c is called the *characteristic leaf* of f_c.

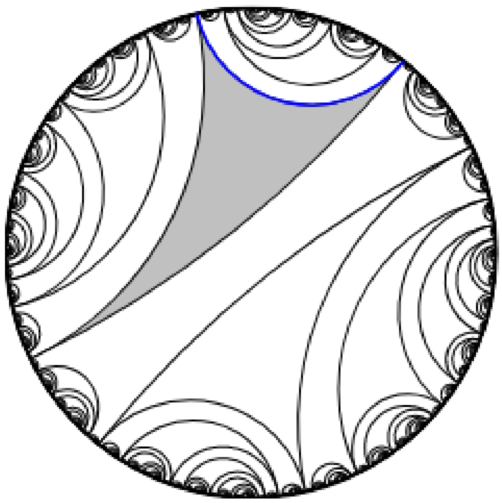

Figure 37: The lamination Λ_ω constructed from the Douady rabbit quadratic f_ω. The characteristic leaf of f_ω is shown in blue. Since f_ω is hyperbolic, it may be shown that $J(\omega)$ is locally connected, so that $\widetilde{K(\omega)}$ is homeomorphic to $K(\omega)$. The grey gap shown above is mapped by the quotient map π_ω to the point $\mu_\omega(\alpha)$, where $\alpha \approx -0.276 + 0.48i$ is the least repelling fixed point of f_ω (see Figure 26 - there, α is shown as the landing point of the 3 rays $R_\omega(1/7), R_\omega(2/7)$ and $R_\omega(4/7)$, which are precisely the rays forming the grey gap shown in this figure). This image is taken from [Exall, 2010], where it is Figure 2.6 on page 21.

Given that $J(c)$ is locally connected, the lamination Λ_c can be used to construct a topological space known as the *abstract filled-in Julia set for f_c* or the *pinched disc model of $K(c)$*. As the name suggests, we begin with the space $D(0,1)$. Define an equivalence relation $\sim_c$ on $D(0,1)$ in the following way. Let $z_1, z_2 \in D(0,1)$ be arbitrary. Of course if $z_1 = z_2$, then we must have $z_1 \sim_c z_2$, so suppose that $z_1 \neq z_2$. If there is a leaf $l \in \Lambda_c$ such that $z_1, z_2 \in l$, then $z_1 \sim_c z_2$. Otherwise, if there is a gap $G \in \mathcal{G}(\Lambda_c)$ such that G has a finite periphery and $z_1, z_2 \in G$, then $z_1 \sim_c z_2$. If no such leaf or gap exists, then $z_1 \not\sim_c z_2$. That is, $\sim_c$ is the smallest equivalence relation on $D(0,1)$ which idenifies all points of a given leaf, and which identifies all points of a given (finite periphery) gap.

Define $\widetilde{K(c)} = D(0,1)/\sim_c$, i.e. $\widetilde{K(c)}$ is the topological space given by taking the disc $D(0,1)$, and then pinching each leaf and each gap with finite periphery into a single point. $\widetilde{K(c)}$ is the pinched disc model of $K(c)$. We note briefly that it is possible to embed $\widetilde{K(c)}$ as a subset of $\mathbb{C}$. Let $\pi_c : D(0,1) \to \widetilde{K(c)}$ be the quotient map.

It is possible to obtain a homeomorphism $\mu_c : K(c) \to \widetilde{K(c)}$. We first define μ_c on the boundary. For any $z \in J(c)$, consider the pre-image $\gamma_c^{-1}(z)$ of z under the parametrisation of $J(c)$. Since only finitely many dynamical rays can land at z, this pre-image is finite, say $\gamma_c^{-1}(z) = \{z_1, \dots, z_k\}$ with $z_i = e^{2\pi i t_i}$ for each $i = 1, \dots, k$. Then the dynamical rays landing at z are $R_c(t_1), \dots, R_c(t_k)$, so that by definition of Λ_c, we have

$$\mathrm{HG}(z_1, z_2), \quad \mathrm{HG}(z_2, z_3), \quad \dots \quad \mathrm{HG}(z_{k-1}, z_k), \quad \mathrm{HG}(z_k, z_1) \quad \in \quad \Lambda_c$$

Then by definition of $\sim_c$, we must have

$$z_1 \sim_c z_2 \sim_c \dots \sim_c z_{k-1} \sim_c z_k$$

and therefore $\pi_c(z_1) = \dots = \pi_c(z_k)$. Then we simply set $\mu_c(z) = \pi_c(z_1)$. That is, a point on the Julia set is mapped to the unique equivalence class of $\widetilde{K(c)}$ that contains the points $e^{2\pi i t}$, where t is the external angle of any one of the dynamical rays landing at z.

Now suppose $z \in K(c) \backslash J(c)$. Then $z \in \mathrm{int}(K(c))$. Let U be the connected component of $\mathrm{int}(K(c))$ that contains z. Since $c \in M$, it may be shown that U is conformally equivalent to $B(0,1)$. Since $J(c)$ is locally connected, there are rays landing at all points of ∂U. These rays will correspond to some collection of leaves in Λ_c which will form a gap whose periphery is infinite. The subset $\pi_c(U) \subseteq \widetilde{K(c)}$ will be conformally equivalent to $B(0,1)$ in the embedding of $\widetilde{K(c)}$ in $\mathbb{C}$. It follows that we may map $\overline{U}$ to $\overline{\pi_c(U)}$ in a homeomorphic way. For all $u \in U$, define $\mu_c(u)$ so that $\mu_c|_{\overline{U}} : \overline{U} \to \overline{\pi_c(U)}$ is some such homeomorphism.

Note that every part of the construction of $\widetilde{K(c)}$ was dependent on the parametrisation $\gamma_c : C(0,1) \to J(c)$ given by the landing of all the rays, which exists if and only if $J(c)$ is locally connected. For $c \in M$ such that $J(c)$ is not locally connected, we are still able to construct a pinched disc model $\widetilde{K(c)}$ for $K(c)$ in a similar (albeit slightly more technical) way. In this case, $\widetilde{K(c)}$ will be locally connected while $K(c)$ will not be locally connected. Thus we certainly don't have a homeomorphism, but we do obtain a continuous surjection $\mu_c : K(c) \to \widetilde{K(c)}$. In this sense, the injectivity of the map μ_c is equivalent to the local connectivity of $J(c)$. As is so often the case, this situation has a nice parallel in the parameter space.

5.3 The Quadratic Minor Lamination and the pinched disc model of M

We were able to construct an abstract topological model for a locally connected Julia set $J(c)$ precisely because of the local connectivity, which tells us that all rays land so that we can construct the lamination Λ_c. Without assuming MLC, we cannot immediately do the same for M, as we do not know that all parameter rays land (recall that all parameter rays with rational external angle do indeed land, and that the landing of all parameter rays with irrational external angles is equivalent to the MLC conjecture). Instead, we must take a more roundabout route. Recall that the minor leaf of a quadratic invariant lamination L is the unique leaf $l \in L$ which is the image of all of the leaves of L that have maximal boundary length.

Definition 5.13 The *Quadratic Minor Lamination* is the set QML consisting of all the leaves that are the minor leaf of some quadratic invariant lamination.

Since every quadratic invariant lamination appears as the lamination Λ_c for some $c \in M$, the quadratic minor lamination is precisely the collection of all characteristic leaves of quadratic maps. The front page of this document displays an image of an approximation of QML, created using *Lavaurs' algorithm*.

Proposition 5.14 *QML is a quadratic invariant lamination.*

Suppose M is locally connected. In this case, the relationship between QML and M is precisely the same as the relationship between Λ_c and $K(c)$. In particular, $\mathrm{HG}(e^{2\pi i t_1}, e^{2\pi i t_2}) \in QML$ if and only if the parameter rays $R_M(t_1)$ and $R_M(t_2)$ land at the same point of ∂M. This does not hold if M is not locally connected, since then we would have some leaf $\mathrm{HG}(e^{2\pi i t_1}, e^{2\pi i t_2}) \in QML$ for irrational t_1, t_2, with the parameter rays $R_M(t_1)$ and $R_M(t_2)$ not landing at all.

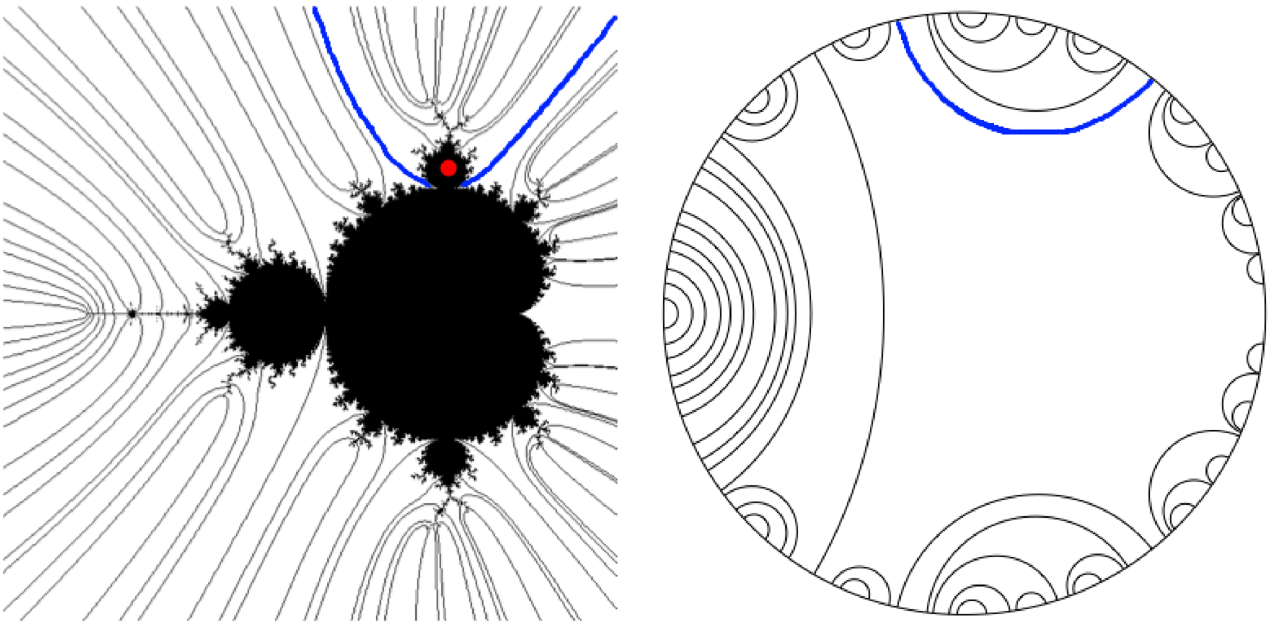

Figure 38: Some paramter rays of M, and the corresponding minor leaves in QML. In QML, the minor leaf of Λ_ω is drawn in blue (see Figure 37). It corresponds to the pair of rays drawn in blue which land at the root of the hyperbolic component containing ω (whose location is marked in red). Any quadratic in this hyperbolic component will have the same lamination as f_ω, and in particular will have their characteristic leaf being the blue leaf above. This image has been modified from the form it takes in its source, which is https://mostlymaths.net, as created by Ruben Berenguel.

QML has two different types of gaps. Recall that the periphery $P(G)$ of a gap G is the intersection of the gap with $C(0,1)$.

Proposition 5.15 *Let $G \in \mathcal{G}(QML)$. Then either:*

(i) $P(G)$ is finite, and we say G is a P-gap, *or*

(ii) $P(G)$ is countably infinite, and we say G is an H-gap.

There are a fair few interesting observations to be made about P-gaps and H-gaps of QML. For example, there is a bijection between P-gaps of QML and Misiurewicz parameters in M, and similarly a bijection between H-gaps and hyperbolic components, with the bijections preserving certain shared features between the gaps and the subsets of M. For more details on this, see any of the 3 excellent references given at the start of this section.

We are now in a position to define an equivalence relation $\sim_M$ on $D(0,1)$. Let $z_1, z_2 \in D(0,1)$ be arbitrary. Of course if $z_1 = z_2$, then we must have $z_1 \sim_M z_2$, so suppose that $z_1 \neq z_2$. If there is a leaf $l \in QML$ such that $z_1, z_2 \in l$, then $z_1 \sim_M z_2$. Otherwise, if there is a P-gap $G \in \mathcal{G}(QML)$ such that $z_1, z_2 \in G$, then $z_1 \sim_M z_2$. If no such leaf or P-gap exists, then $z_1 \not\sim_M z_2$. That is $\sim_M$ is the smallest equivalence relation on $D(0,1)$ which identifies all points on a given leaf, and which identifies all points on a given P-gap.

Just as in the dynamical case, we obtain the pinched disc model $\widetilde{M} = D(0,1)/\sim_M$ of M, by pinching each leaf and each P-gap of QML into a single point, and the space $\widetilde{M}$ can be embedded as a subset of $\mathbb{C}$. The below theorem from [Douady, 1993] is also in parallel to the dynamical case.

Theorem 5.16 *There exists a continuous surjection $\mu_M : M \to \widetilde{M}$ which is injective if and only if M is locally connected. Thus, M is homeomorphic to $\widetilde{M}$ if and only if M is locally connected.*

The "only if" part of the latter statement follows because it may be shown that $\widetilde{M}$ is locally connected.

Let $\pi_M : D(0,1) \to \widetilde{M}$ be the quotient map. The map μ_M is constructed in the following way. For a hyperbolic component H, there is a unique H-gap G of QML corresponding to H. In the embedding of $\widetilde{M}$ in the plane, we will have that $\pi_M(G)$ is conformally equivalent to $B(0,1)$. It follows that $\overline{H}$ and $\overline{\pi_M(G)}$ can be mapped to one another in some homeomorphic way, and we let $\mu_M|_{\overline{H}}$ be some such way. For other points, we essentially consider a sequence of rational parameter ray pairs (two rays with rational external angle that land at a shared point) whose wakes form a nested sequence of subsets, all containing the point. These wakes become as small as is possible so that they still contain the point. This gives a sequence of leaves in QML that will have a limit leaf. Then, π_M will identify the entire limit leaf into one point in $\widetilde{M}$, which is where we map our point from M. This is a rather informal and perhaps not wholly accurate description of the process, which is best described using *fibers* of M - see [Schleicher, 2009] for the full details.

5.4 The implication (MLC) $\Rightarrow$ (DHC)

Just as in the classical text [Douady and Hubbard, 1985a], our final section is dedicated to a proof of Theorem 3.18, i.e. that local connectivity of M implies density of hyperbolicity for complex quadratic maps! Unlike Douady and Hubbard, our proof is the roughest of sketches, and is as presented in [Schleicher, 2009]. Recall that a *ghost component* U of M is a connected component of $\text{int}(M)$ such that for all $c \in U$, the map f_c has no attracting cycles. Density of hyperbolicity for complex quadratic maps is equivalent to the non-existence of ghost components in M (see the concluding remarks of Section 3.2).

It is via the map μ_M that we show the implication. Along with the information about μ_M given in Theorem 5.16, we will require the following, which follows from the construction of μ_M in [Schleicher, 2009], as well as the following.

Proposition 5.17 *Suppose U is a ghost component of M. Then μ_M is constant on U, i.e the image of any given ghost component of M is a single point of $\widetilde{M}$.*

This is a consequence of the fact that ghost components are 'contained in a single fiber', itself a consequence of Douady and Hubbard's *Branch Theorem* which we mention precisely because of its significance in showing that MLC implies DHC. Proposition 5.17 allows us to bring in density of hyperbolicity.

Corollary 5.18 *Suppose $\text{int}(\mu_M^{-1}(x)) = \varnothing$ for all $x \in \widetilde{M}$. Then hyperbolic maps are dense in the space of complex quadratic maps.*

Proof. We instead demonstrate the contrapositive statement:

If hyperbolicity is not dense for complex quadratics, then there exists $x \in \widetilde{M}$ such that $\text{int}(\mu_M^{-1}(x)) \neq \varnothing$.

Suppose hyperbolicity is not dense for complex quadratics. Then M has some ghost component U. Since U is a connected component of the interior of M, it has a non-empty interior. By Proposition 5.17, μ_M is constant on U, say $\mu_M(U) = \{x\}$ for some $x \in \widetilde{M}$. Then $\mu_M^{-1}(x) = U$ has a non-empty interior, as required. □

The converse is also true, and would quickly follow, but we have no use for it. Now, we may finish with:

Proof of Theorem 3.18: Suppose M is locally connected. Then by Theorem 5.16, μ_M is a bijection. Let $x \in \widetilde{M}$ be arbitrary. Then the pre-image of x is $\mu_M^{-1}(x) = \{c\}$, for some $c \in M$. In particular, $\text{int}(\mu_M^{-1}(x)) = \varnothing$, so by Corollary 5.18, hyperbolic maps are dense in the space of complex quadratic maps.

References

[Beardon, 2000] Beardon, A. (2000). *Iteration of Rational Functions: Complex Analytic Dynamical Systems*. Graduate Texts in Mathematics. Springer New York.

[Benini, 2018] Benini, A. M. (2018). A survey on MLC, Rigidity and related topics.

[Blokh et al., 2017] Blokh, A., Oversteegen, L., and Timorin, V. (2017). Perfect subspaces of quadratic laminations.

[Branner and Fagella, 2001] Branner, B. and Fagella, N. (2001). Extensions of homeomorphisms between limbs of the Mandelbrot set. *Conformal Geometry and Dynamics of The American Mathematical Society*, 5:100–139.

[Branner and Hubbard, 1988] Branner, B. and Hubbard, J. H. (1988). The iteration of cubic polynomials Part I: The global topology of parameter space. *Acta Math.*, 160:143–206.

[Carleson and Gamelin, 1996] Carleson, L. and Gamelin, T. (1996). *Complex Dynamics*. Universitext. Springer New York.

[Cheraghi and Shishikura, 2015] Cheraghi, D. and Shishikura, M. (2015). Satellite renormalization of quadratic polynomials.

[Devaney and Keen, 1997] Devaney, R. L. and Keen, L. (1997). *Chaos and Fractals: The Mathematics behind the Computer Graphics*. American Mathematical Society, USA, 1st edition.

[Douady, 1993] Douady, A. (1993). Descriptions of compact sets in $\mathbb{C}$. In *Topological Methods in Modern Mathematics: A Symposium in Honor of John Milnor's Sixtieth Birthday*, pages 429–465. Publish or Perish, Inc.

[Douady and Hubbard, 1985a] Douady, A. and Hubbard, J. H. (1984-1985a). Exploring the Mandelbrot set. The Orsay Notes.

[Douady and Hubbard, 1985b] Douady, A. and Hubbard, J. H. (1985b). On the dynamics of polynomial-like mappings. *Annales scientifiques de l'École Normale Supérieure*, 18(2):287–343.

[Dudko and Sutherland, 2020] Dudko, A. and Sutherland, S. (2020). On the Lebesgue measure of the Feigenbaum Julia set. *Inventiones mathematicae*, 221(1):167–202.

[Dudko and Lyubich, 2019] Dudko, D. and Lyubich, M. (2019). Local connectivity of the Mandelbrot set at some satellite parameters of bounded type.

[Epstein and Yampolsky, 1999] Epstein, A. and Yampolsky, M. (1999). Geography of the cubic connectedness locus: Intertwining surgery. *Annales Scientifiques de l'École Normale Supérieure*, 32(2):151 – 185.

[Exall, 2010] Exall, F. (2010). *Rational maps represented by both rabbit and aeroplane matings*. PhD thesis, University of Liverpool.

[Fatou, 1919] Fatou, P. (1919). Sur les équations fonctionnelles. *Bulletin de la Société mathématique de France*, 47:161–271.

[Graczyk and Świątek, 1997] Graczyk, J. and Świątek, G. (1997). Generic Hyperbolicity in the Logistic Family. *Annals of Mathematics*, 146.

[Graczyk and Świątek, 1998] Graczyk, J. and Świątek, G. (1998). *The real Fatou conjecture*. Annals of Mathematics Studies, 144. Princeton University Press, Princeton, NJ.

[Hu and Jiang, 1993] Hu, J. and Jiang, Y. (1993). The Julia set of the Feigenbaum quadratic polynomial is locally connected.

[Hubbard, 1993] Hubbard, J. H. (1993). Local connectivity of Julia sets and bifurcation loci: three theorems of J.-C. Yoccoz. In *Topological methods in modern mathematics (Stony Brook, NY, 1991)*.

[Inou and Shishikura, 2008] Inou, H. and Shishikura, M. (2008). The renormalization for parabolic fixed points and their perturbations.

[Jiang, 1995] Jiang, Y. (1995). Local connectivity of the Mandelbrot set at certain infinitely renormalizable points.

[Julia, 1918] Julia, G. (1918). Mémoire sur l'itération des fonctions rationnelles. *Journal de Mathématiques Pures et Appliquées*, 1:47–246.

[Kahn, 2006] Kahn, J. (2006). A priori bounds for some infinitely renormalizable quadratics: I. Bounded primitive combinatorics.

[Kahn and Lyubich, 2006] Kahn, J. and Lyubich, M. (2006). A priori bounds for some infinitely renormalizable quadratics: II. Decorations.

[Kahn and Lyubich, 2008] Kahn, J. and Lyubich, M. (2008). A priori bounds for some infinitely renormalizable quadratics: III. Molecules.

[Kozlovski et al., 2007] Kozlovski, O., Shen, W., and Strien, S. (2007). Density of hyperbolicity in dimension one. *Annals of Mathematics*, 166.

[Lavaurs, 1989] Lavaurs, P. (1989). *Systèmes dynamiques holomorphes : explosion de points périodiques paraboliques*. PhD thesis, Université Paris-Sud 11. Thèse de doctorat dirigée par Douady, Adrien Mathématiques Paris 11 1989.

[Lei, 1990] Lei, T. (1990). Similarity between the Mandelbrot set and Julia sets. *Comm. Math. Phys.*, 134(3):587–617.

[Lei, 1992] Lei, T. (1992). Voisinages connexes des points de Misiurewicz. *Annales de l'institut Fourier*, 42(4):707–735.

[Lei, 2000] Lei, T. (2000). Local properties of the Mandelbrot set at parabolic points. In *The Mandelbrot set: Theme and Variations*, volume 274 of *London Math. Soc. Lecture Note Series*, pages 133–160. Cambridge Univ. Press.

[Levin, 2007] Levin, G. (2007). Rigidity and non local connectivity of Julia sets of some quadratic polynomials.

[Levin, 2011] Levin, G. (2011). Addendum to: "Rigidity and non local connectivity of Julia sets of some quadratic polynomials". *Communications in Mathematical Physics*, 325:1171–1178.

[Lyubich, 1997] Lyubich, M. (1997). Dynamics of quadratic polynomials, I–II. *Acta Mathematica*, 178(2):185 – 297.

[Mandelbrot, 1985] Mandelbrot, B. B. (1985). On the dynamics of iterated maps. V. Conjecture that the boundary of the M-set has a fractal dimension equal to 2. In *Chaos, fractals, and dynamics (Guelph, Ont., 1981/1983)*, Lecture Notes in Pure and Appl. Math., 98.

[McMullen, 1994a] McMullen, C. T. (1994a). *Complex Dynamics and Renormalization (AM-135)*. Princeton University Press.

[McMullen, 1994b] McMullen, C. T. (1994b). Frontiers in complex dynamics. *Bulletin of the American Mathematical Society*, 31(2):155–173.

[Milnor, 1990] Milnor, J. (1990). Dynamics in one complex variable: introductory lectures.

[Milnor, 1992] Milnor, J. (1992). Remarks on Iterated Cubic Maps. *Exp. Math.*, 1:5–24.

[Peitgen et al., 1992] Peitgen, H., Jurgens, H., and Saupe, D. (1992). *Chaos and Fractals (New Frontiers of Science)*, volume 1. Springer.

[Schleicher, 1997] Schleicher, D. (1997). Rational parameter rays of the Mandelbrot set. *arXiv Mathematics e-prints*, page math/9711213.

[Schleicher, 1999] Schleicher, D. (1999). On Fibers and Local Connectivity of Mandelbrot and Multibrot Sets.

[Schleicher, 2009] Schleicher, D. (2009). Laminatins, Julia sets, and the Mandelbrot set; appendix to "The combinatorics of iterated rational maps". In *Complex Dynamics: Families and Friends (2009)*, pages 111–130. Taylor & Francis.

[Shishikura, 1998] Shishikura, M. (1998). The Hausdorff dimension of the boundary of the Mandelbrot set and Julia sets. *Annals of Mathematics*, 147:225–267.

[Sullivan, 1988] Sullivan, D. (1988). Bounds, quadratic differentials, and renormalization conjectures. Technical Report IHES-M-91-25, Inst. Hautes Etud. Sci., Bures-sur-Yvette.

[Thurston, 1985] Thurston, W. (1985). The combinatorics of iterated rational maps. In *Complex Dynamics: Families and Friends (2009)*, pages 1–108. Taylor & Francis.

[Torhorst, 1921] Torhorst, M. (1921). *Über den Rand der einfach zusammenhängenden ebenen Gebiete.* PhD thesis, Rheinische Friedrich-Wilhelms-Universität Bonn.