

How to prune a horseshoe

André de Carvalho* and Toby Hall†

Abstract

Let $F: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ be a homeomorphism. An open F -invariant subset U of \mathbb{R}^2 is a *pruning region* for F if it is possible to deform F continuously to a homeomorphism F_U for which every point of U is wandering, but which has the same dynamics as F outside of U . This concept is motivated by the *Pruning Front Conjecture* introduced by Cvitanović which claims that every Hénon map can be understood as a pruned horseshoe.

This paper contains recent results on pruning theory, concentrating on prunings of the horseshoe. We give conditions on a disk D which ensure that the orbit of its interior is a pruning region; explain how prunings of the horseshoe can be understood in terms of underlying tree maps; discuss the connection between pruning and Thurston's classification theorem for surface homeomorphisms; motivate a conjecture describing the *forcing relation* on horseshoe braid types; and use this theory to give a precise statement of the Pruning Front Conjecture.

1 Introduction

This paper describes some recent developments in pruning theory. It is aimed at both mathematicians and physicists with an interest in the mechanism of transitions to chaos in parameterized families such as the Hénon family. Although the basic ideas of the theory are simple and geometric, the calculations involved in studying even relatively straightforward examples can act as a deterrent, descending as they do either into a morass of symbolic computations, or into complicated diagrams. It is hoped that these difficulties will be overcome at least in part by the example-based approach adopted here.

The term 'pruning' was coined by Cvitanović [26] in a paper which presented a model for the topological dynamics of the Hénon family: each Hénon map is understood as a horseshoe from which some of the dynamics has been 'pruned away'. Using this model Cvitanović, Gunaratne, and Procaccia [28] were able to predict with great accuracy quantities such as the number of periodic orbits of small (< 16) period and the topological entropy of maps in the Hénon and Lozi families. More importantly, pruning provided a new approach to the topological description of families of plane homeomorphisms, analogous to the kneading theory of Milnor and Thurston. The deep ideas introduced by Cvitanović were developed by the first author of the present paper, who introduced a mathematical theory of pruning as a mechanism for destroying the dynamics of surface homeomorphisms in a controlled way. Since then, pruning has been used to give an alternative

*Institute for Mathematical Sciences, State University of New York, Stony Brook, NY 11794-3660, USA, andre@math.sunysb.edu

†Department of Mathematical Sciences, University of Liverpool, Liverpool L69 7ZL, UK, T.Hall@liverpool.ac.uk

constructive approach to Thurston’s classification of surface homeomorphisms, and to motivate a conjecture about the way in which Boyland’s forcing relation organizes the set of horseshoe braids. It can also be used to give a precise statement of the original Pruning Front Conjecture.

This paper contains a description of work motivated by [26, 28], bringing together results from several of the authors’ papers as well as previously unpublished work. The main ideas described are the following:

- *Pruning disks and the pruning theorem:* Pruning disks are the main tool for constructing models for Hénon and Lozi maps starting with Smale’s horseshoe. A pruning disk D for a homeomorphism F (such as the horseshoe) is a disk with the property that it is possible to destroy all of the orbits of F which enter D , while leaving other orbits untouched. The pruning theorem asserts that if D satisfies certain *pruning conditions*, then F can be deformed continuously to a homeomorphism which is equal to F at every point whose orbit doesn’t enter D , but for which every point of D is wandering.
- *Pruning and tree maps:* Pruned horseshoes can be understood in terms of tree maps. This makes it possible to find Markov partitions for pruned horseshoes (or good approximations, if there is no Markov partition); to describe the symbolics of orbits which survive the pruning in terms of the original symbolic dynamics of the horseshoe; and hence to compute dynamical ζ -functions.
- *Maximal prunings:* Using the relationship between pruned maps and tree maps, an algorithm is developed for finding a maximal pruning of the horseshoe which avoids (or *is relative to*) given periodic or homoclinic orbits. This provides an approach to the problem of orbit forcing in the horseshoe: the orbits which survive a maximal pruning relative to a given orbit P are precisely those which are forced by P .
- *Families of orbits:* The calculation of maximal prunings motivates a conjecture which organizes all of the periodic orbits of the horseshoe into families which are linearly ordered by forcing. The families themselves are parameterized by orbits homoclinic to the outer fixed point of the horseshoe. The linear order within families is simply the unimodal order on the symbolic codes of the orbits. Some special cases of the conjecture have been proved. Before the necessary preliminaries have been covered, it is hard to give examples which properly reflect the scope of this conjecture — as a very special case, statements such as ‘if a Hénon map has an orbit with symbolic representation $\overline{1000100}$, then it also has orbits with symbolic representations $\overline{100(1100)^n100}$ for all $n \geq 1$ ’ follow straightforwardly from it. (Although the conjecture is unproved, the above example is covered by one of the special cases which has been proved.)
- *Braid types and forcing:* The *braid type* of a periodic orbit of a plane homeomorphism is an analogue of the permutation of a periodic orbit in dimension 1. For unimodal maps of the line, the permutation of a periodic orbit and its symbolic representation contain essentially the same information, but in dimension 2 this is not the case: orbits with quite different symbolics can have the same braid type. Braid types are therefore important in the study of the Hénon family, where the absence of critical points makes it difficult to introduce symbolic dynamics unambiguously. This makes it possible to build a framework for the proof of the Pruning Front Conjecture which doesn’t rely on symbolic dynamics as a starting point. The

linear order on the families discussed above is also formalized in terms of braid type forcing, which is defined in the same way as permutation forcing in dimension 1.

- *The Pruning Front Conjecture (PFC)*: The Pruning Front Conjecture as introduced by Cvitanović was described intuitively, and relied on numerical evidence to validate the models. In order to give a precise statement of it, a *0-entropy equivalence relation* is introduced: this identifies points which can be joined by a compact connected set carrying no entropy. The *Braid Type Conjecture (BTC)* is an analogue of PFC which states only that the periodic orbits of Hénon maps are of the same topological type as periodic orbits of the horseshoe. It is explained how a proof of BTC might provide an approach to a proof of PFC.

At several points in the paper we have chosen clarity at the expense of mathematical precision, although on each such occasion there is an exculpatory remark either in the text or in a footnote. The ultimate goal of all of the ideas described here is a complete description of the possible topological dynamics presented by parameterized families of maps which pass from trivial to chaotic dynamics (what we mean by ‘complete’ is explained in Section 7). In particular, we have been motivated by working towards a proof of the Pruning Front Conjecture. Ishii [63] proved the Pruning Front Conjecture for the Lozi family, but his approach is adapted to the piecewise linear nature of this family and is quite different from the one adopted here. There is a vast literature on the Hénon and Lozi families, the transition to chaos in dimension 2, and pruning theory in other contexts (see for example¹ [18, 17, 60, 61, 10, 9, 11, 12, 14, 16, 13, 15, 8, 7, 78, 53, 54, 52, 56, 63, 64, 29, 77, 65, 26, 27, 25, 2, 3, 46, 45]). Most of this work is not discussed here: the topological approach described in this paper has a quite different flavour.

After some necessary preliminaries have been dealt with in Section 2, we introduce the concept of pruning disks and pruning fronts in Section 3, culminating with a statement of the pruning theorem of [32]. Section 4 contains the bulk of the theory which is applied in the remainder of the paper. It describes how pruning plane homeomorphisms can be interpreted in terms of a generalized kneading theory for tree maps, discusses the concept of efficient tree maps as introduced by Bestvina and Handel, and indicates the main steps in the pruning approach to Thurston’s classification theorem for surface homeomorphisms. Section 5 is concerned with families of efficient tree maps, and contains the motivation for a conjecture about the orders in which periodic orbits can be formed in the creation of a horseshoe. In Section 6 Boyland’s concept of braid types is introduced into the discussion, and the relationship between Thurston’s classification theorem, pruning, and forcing is explained. The Pruning Front Conjecture is stated in Section 7, and an approach to its proof (using another conjecture, the Braid Type Conjecture) is outlined. Finally, Section 8 summarizes the paper and describes some open problems. Two appendices give detailed examples of computational aspects of the theory: Appendix A shows how to construct the tree map corresponding to a given pruning front, and Appendix B contains an example of how to use the tree map to determine the (horseshoe) symbolics of orbits which survive pruning.

The confusion which often seems to be generated by many of the diagrams in this paper can be lessened by using the program [47] which was used to produce them. The reader is strongly encouraged to use this program to work through the examples presented here and other similar ones.

Acknowledgements: We would like to thank P. Cvitanović and S. Peles for their interest in this

¹This is certainly not an exhaustive list of papers on these subjects.

work, and for their many helpful comments about the exposition of this paper, particularly those aimed at making it more palatable to physicists. The comments of the three referees have also been very valuable. We are grateful to the Institute for Mathematical Sciences at SUNY, Stony Brook, which has provided financial support for much of the work presented here, as well as a friendly and stimulating work atmosphere. The second author is also grateful for the financial support provided by the Clay Institute for a visit to the USA.

2 The horseshoe and the Hénon family

Smale's horseshoe diffeomorphism [74] is an excellent example of a 2-dimensional dynamical system having very rich dynamics which can be completely understood. The following brief description is intended primarily to introduce the notation that will be used throughout the paper: good detailed treatments can be found in many texts on dynamical systems (e.g. [38]). The horseshoe is a diffeomorphism F of the sphere (or of the plane), all of whose interesting dynamics lies inside a stadium shaped region denoted by \mathbb{T} in Figure 1: \mathbb{T} is the union of a square $S = (0, 1) \times [0, 1]$, and two half-disks of diameter 1 centred on $(0, 1/2)$ and $(1, 1/2)$. The action of $F|_{\mathbb{T}}$ is to stretch S uniformly in the horizontal direction by a factor λ (which must be greater than 2), to contract S uniformly in the vertical direction by a factor $1/\lambda$, to bend the deformed S into the shape of a horseshoe, and to map it into \mathbb{T} as shown in the figure, in such a way that vertical segments in $S \cap F^{-1}(S)$ map into vertical segments in S , and horizontal segments in $S \cap F^{-1}(S)$ map onto horizontal segments in S . The two half-disks on the left and right are contracted uniformly and mapped inside the one on the left. Outside of \mathbb{T} , F has a fixed point at ∞ with $F^{-n}(x) \rightarrow \infty$ as $n \rightarrow \infty$ for all $x \in S^2 \setminus \mathbb{T}$.

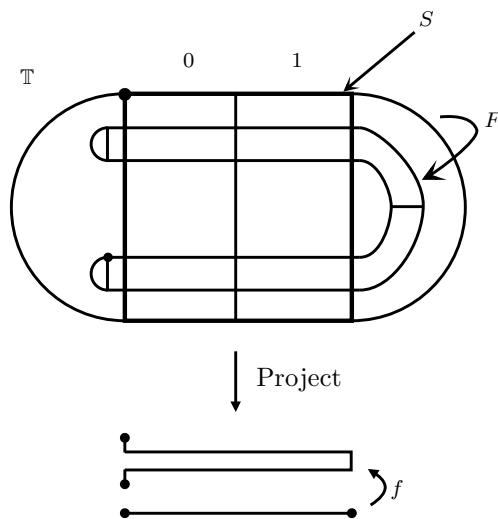


Figure 1: The horseshoe, its symbolic dynamics, and its one-dimensional quotient

The non-wandering set $\Omega(F)$ of F is the union of ∞ , an attracting fixed point inside the left half-disk and a Cantor set

$$\Lambda = \{x \in S : F^n(x) \in S \text{ for all } n \in \mathbb{Z}\}.$$

A fact that will be of central importance in this paper is that the horseshoe projects to a self-map of the interval. Indeed, from the way we chose the standard model for the horseshoe, it is clear that collapsing to points a) each vertical segment in S , and b) each of the two half-disks yields the interval $I = [0, 1]$: and $F: \mathbb{T} \rightarrow \mathbb{T}$ projects to the *full tent map* $f: I \rightarrow I$ whose image covers I twice as shown in Figure 1.

The dynamics of F on Λ can be understood symbolically: let $p: \mathbb{T} \rightarrow [0, 1]$ be the projection just described, and define the *itinerary map* $k: \Lambda \rightarrow \Sigma = \{0, 1\}^{\mathbb{Z}}$ by

$$k(x)_i = \begin{cases} 0 & \text{if } p(F^i(x)) < 1/2 \\ 1 & \text{if } p(F^i(x)) > 1/2. \end{cases}$$

Then k is a homeomorphism which conjugates the action of F on Λ to that of the *shift map* $\sigma: \Sigma \rightarrow \Sigma$, defined by $\sigma(s)_i = s_{i+1}$: that is, $F|_{\Lambda} = k^{-1} \circ \sigma \circ k$. Thus every point x of Λ has a symbolic representation $k(x) \in \Sigma$, and it is often convenient to refer to x by this symbolic representation without explicitly invoking the homeomorphism k . In particular, periodic points of F correspond one-to-one with periodic sequences of the same period. When writing elements $s \in \Sigma$, it is common to juxtapose a point between s_{-1} and s_0 to indicate the origin of the sequence, and to use an overbar to denote infinite repetition of a given word either at the beginning or the end of s . When s is periodic, and hence can be written in the form \overline{w} for some word w , then the origin of the sequence is taken to be at the first letter of w , i.e. $\overline{w} = \dots www \cdot www \dots$.

The horizontal and vertical ordering of points of Λ can be determined using the *unimodal order* \preceq on $\Sigma_+ = \{0, 1\}^{\mathbb{N}}$, which is defined as follows: if $s = s_0s_1\dots$ and $t = t_0t_1\dots$ are distinct elements of Σ_+ , and $n \in \mathbb{N}$ is least such that $s_n \neq t_n$, then $s \prec t$ if and only if $\sum_{i=0}^n s_i$ is even. If $s, t \in \Sigma_+$, then $s \preceq t$ if and only if either $s = t$ or $s \prec t$. Now if $(x_1, y_1), (x_2, y_2) \in \Lambda$ with $k(x_1, y_1) = \dots s_{-2}s_{-1} \cdot s_0s_1\dots$ and $k(x_2, y_2) = \dots t_{-2}t_{-1} \cdot t_0t_1\dots$, then

$$\begin{aligned} x_1 < x_2 &\iff s_0s_1s_2\dots \prec t_0t_1t_2\dots \text{ and} \\ y_1 < y_2 &\iff s_{-1}s_{-2}\dots \prec t_{-1}t_{-2}\dots \end{aligned}$$

Figure 2 depicts the symbol square (see the Notational Interlude in Section 4.2 for further clarification), together with the partial addresses of the sixteen boxes obtained using two forward and two backward iterations of the horseshoe.

01-00	01-01	01-11	01-10
11-00	11-01	11-11	11-10
10-00	10-01	10-11	10-10
00-00	00-01	00-11	00-10

Figure 2: The symbol square

Notice that vertical segments of S which contain elements of Λ can be specified by an element of Σ_+ .

A period n orbit P of F will be described by its *code* $c_P \in \{0, 1\}^n$, which is given by the first n symbols of the itinerary of its rightmost point p : thus $k(p) = \overline{c_P}$. For example, the period 5 orbit which contains the point with itinerary $\overline{01001}$ has code 10010. Thus a word $c \in \{0, 1\}^n$ is the code of some period n horseshoe orbit if and only if the semi-infinite sequence \bar{c} is strictly greater than $\sigma^i(\bar{c})$ in the unimodal order for $1 \leq i < n$. This paper is also concerned with homoclinic orbits H to the fixed point with code 0: in the remainder of the paper, ‘homoclinic’ will always mean homoclinic to this fixed point.

The dynamics of the tent map f can be modelled by the 1-sided 2-shift, and the quotient described above induces a 1-1 correspondence between the periodic orbits of f and those of F (two periodic orbits correspond if and only if they have the same code). We will often speak about periodic orbits of both maps as though they were the same without mentioning this identification.

The Hénon and Lozi families will be important examples throughout the paper. The first [57] is the family of diffeomorphisms of the plane defined by

$$H_{a,b}(x, y) = (a - x^2 - by, x),$$

and the second [70] is the family of piecewise linear plane homeomorphisms defined by

$$L_{a,b}(x, y) = (1 - a|x| - by, x),$$

where a and b are real parameters. It will always be assumed that $b > 0$ (so that the maps in both families are orientation-preserving).

It can be shown [39] that if the parameter a is chosen large enough, then the recurrent dynamics of the Hénon map $H_{a,b}$ is precisely the recurrent dynamics of the horseshoe. For small values of a , on the other hand, it is easy to show that $H_{a,b}$ has trivial dynamics. The question which has engaged both mathematics and physics communities for many years is: *what exactly happens in between?*

Figure 3 shows the stable and unstable manifolds of the outer fixed point (i.e. the one whose eigenvalues are always positive) of the Hénon map $H_{a,b}$ for three choices of the parameters. In a) (where $a = -0.3$, $b = 0.5$) there is no homoclinic intersection, and no interesting dynamics apart from the two fixed points. In c) (where $a = 5.5$, $b = 0.8$) we see the intersection pattern of the stable and unstable manifolds of the fixed point $\bar{0}$ of the horseshoe. In b) (where $a = 1.26$, $b = 0.8$), the situation is clearly more complicated.




Figure 3: Stable and unstable manifolds in the Hénon family

3 The pruning theorem

In this section we state the pruning theorem [32], which describes how it is possible to destroy the dynamics of a surface homeomorphism in a controlled way by a continuous deformation through homeomorphisms (i.e. by an *isotopy*). Although the theorem holds in a more general setting, we will only deal here with homeomorphisms of \mathbb{R}^2 .

Let F be a homeomorphism of the plane (for example the horseshoe, or a Hénon map). Roughly speaking, if F is a diffeomorphism, then a *pruning disk* D for F is a region which i) is bounded by a segment of stable manifold and a segment of unstable manifold intersecting only at their endpoints; and ii) satisfies an additional dynamical condition. The purpose of the additional condition is to ensure that it is possible to deform F so as to remove all crossings of stable and unstable manifolds contained in D , as illustrated in Figure 4. It is not the case that any disk bounded by segments of stable and unstable manifolds serves this purpose. In Figure 5, a disk D and an iterate $F^n(D)$ are shown. If it were possible to uncross the stable and unstable manifolds inside D , then the same would happen inside all of its images and preimages, including $F^n(D)$. The right hand side of Figure 5, however, illustrates that it is impossible to remove the circled crossings.




Figure 4: Uncrossing stable and unstable manifolds inside a pruning disk




Figure 5: A disk which is not a pruning disk

This example suggests, at least on an intuitive level, that if some forward iterate $F^n(C)$ of the stable side C of a disk D intersects the interiors of both D and its complement, then it is impossible to deform F so as to remove all crossings of stable and unstable manifolds in D . On the other hand, Figure 6 illustrates that if $F^n(C)$ is contained entirely in D then uncrossing the manifolds may be possible, but that in this case the same effect could be achieved by replacing D with a larger disk D' with the property that $F^n(C')$ is disjoint from the interior of D' . To

see why this is so, notice that the circled intersections immediately to the left of C are mapped by F^n to intersections immediately to the right of $F^n(C)$, which are contained in D ; hence if all the intersections in D are to be removed, then so must be those to the left of D . Thus D can be enlarged without modifying the set of intersections to be removed. This enlargement can be continued until $F^n(C')$ is contained in C' (and is therefore disjoint from the interior of D'): note that $F^n(C')$ then contains a fixed point p of F^n .

Thus if $F^n(C)$ is contained in the interior of D , then D can be replaced with an equivalent disk D' , with $F^n(C')$ disjoint from the interior of D' . Because of this, the definition of pruning disks given below forbids all intersections of $F^n(C)$ with the interior of D , not just the type depicted in Figure 5.

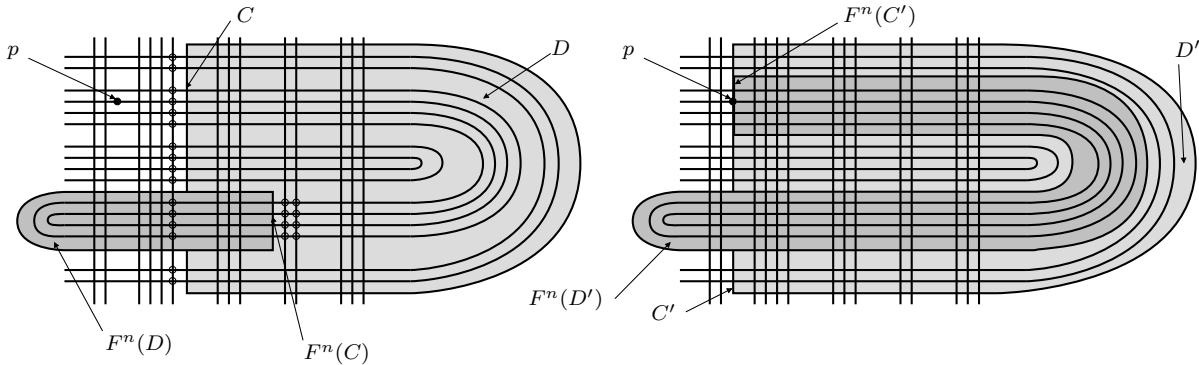


Figure 6: Equivalent pruning disks

Because the pruning theorem constructs an isotopy which need not preserve differentiability², the following definition of pruning disks is made in such a way as not to depend on the existence of stable and unstable manifolds. The arcs C and E in the definition should nevertheless be regarded intuitively as segments of stable and unstable manifold respectively.

A *pruning disk* for F is a closed topological disk D whose boundary can be written as the union of two arcs intersecting only at their endpoints, $\partial D = C \cup E$, satisfying the following *pruning conditions*:

- i) C (respectively E) is contracted by forward (backward) iteration under F . More precisely, its diameter tends to 0 under forward (backward) iteration.
- ii) Forward (backward) iterates of C (E) do not intersect the interior of the disk D .

Remark: When dealing with non-differentiable maps in general, condition ii) above is rather more technical. The version given here is adequate for our purposes in this paper. The interested reader should see [32] for more details.

²The differentiability loving reader should not be discouraged by this. The purpose of the pruning theorem is to construct *models* which describe the topological dynamics of interesting differentiable maps, such as Hénon maps. In Section 7 we describe how to relate the models to the actual maps they are supposed to describe.

Example 1 Figure 7 depicts two pruning disks for the horseshoe. In a), the pruning disk has its E - and C -sides contained in the unstable and stable manifolds of the fixed point of code 0. In b), the E - and C -sides are contained in the unstable and stable manifolds of the other fixed point, with code 1. In each case, pruning condition i) is immediate since the sides of the disks are contained in stable and unstable manifolds. To see that pruning condition ii) holds for a), observe that the forward iterates of C (shown with thicker lines) are all disjoint from D (since they are attracted to the fixed point of code 0, only finitely many need to be considered); a similar simple argument holds for the backward iterates of E . In b), all of the iterates $F^n(C)$ with $n \geq 0$ are contained in the vertical segment through the fixed point, and all of the iterates $F^n(E)$ with $n < 0$ are contained in the horizontal segment through the fixed point: hence none of them intersects the interior of D .

The argument used in b) generalizes readily: if a disk D is bounded by a segment C of the stable manifold and a segment E of the unstable manifold of a fixed point p ; and if neither the segment of stable manifold between p and C nor the segment of unstable manifold between p and E intersects the interior of D , then D is a pruning disk.




Figure 7: Pruning disks for the horseshoe

Given a pruning disk D for a homeomorphism F , the pruning theorem states that it is possible to destroy the dynamics inside the pruning disk (and all of its iterates), leaving the homeomorphism unchanged elsewhere. That is, the orbits which are destroyed are precisely those which enter D . If the resulting map is a diffeomorphism (which need not always be the case), it follows that at the end of the isotopy no intersection of stable and unstable manifolds is left inside D .

Theorem 3.1 *Given a plane homeomorphism F and a pruning disk D for F , there exists an isotopy supported on the interior of D , together with all of its forward and backward iterates, such that, if F_D denotes the map at the end of the isotopy, all points in the interior of D are wandering points of F_D .*

Thus in each of the examples of Figure 7, it is possible to deform the horseshoe so as to destroy

all nonwandering orbits (and in particular periodic orbits) which enter the (interior of the) pruning disk, while leaving the dynamics unchanged elsewhere.³

Remarks:

- a) The map F_D obtained at the end of the isotopy in the pruning theorem above will be referred to variously as *a pruning of F* , *the pruned map*, *the map obtained from F by pruning D away*, etc.
- b) A suitably generalized version of the pruning theorem holds for finite unions of pruning disks, which we refer to as *pruning fronts* (see [32] for details). In the language of Cvitanović [26], it would be more appropriate to call them *primary pruned regions*.
- c) Since F is a homeomorphism, images and preimages of a pruning disk are also pruning disks, and the pruning theorem applied to any of them yields the same result. This remark will be used later to enable us to restrict attention to pruning disks in the horseshoe which extend up to the right hand edge of the square S and are symmetric about its horizontal centre line (such as those of Figure 7).

4 Pruning and tree maps

This section is the longest in the paper and is divided into six subsections. In Section 4.1 it is shown how to use the pruning theorem to construct a family of two-dimensional homeomorphisms analogous to a full family of unimodal interval maps. Section 4.2 describes how pruning is related to a kind of generalized kneading theory for tree maps: pruned maps project to tree maps, and the pruning isotopy corresponds to ‘pulling away’ backtracking pieces of the quotient tree map. This makes it possible to obtain Markov partitions for certain pruned maps, and to construct Markov graphs which describe the symbolics of the surviving orbits in terms of the original 0-1 symbolics of the horseshoe. Details of the construction of tree maps from pruning fronts, and of the Markov graphs which describe surviving orbits, can be found in Appendices A and B respectively. As an example, in Section 4.3 it is shown how the dynamics of the Hénon map with parameters $a = 5.4$ and $b = 1$ may conjecturally be understood in terms of a pruned horseshoe and the associated tree map. This approach is compared with the work of Davis, MacKay, and Sannami [30] on the same Hénon map. In Section 4.4 the concept of *efficient* tree maps is explained. An efficient tree map is one possible output of the Bestvina-Handel algorithm [19] (see also [43, 69]) for finding the Thurston representative in the isotopy class of a surface homeomorphism (see Section 6.1 for more details). This concept will be central in Section 5 when families of periodic orbits are discussed. In the final two subsections we relate pruning and efficiency and explain how to find Thurston representatives using pruning — that is, how to find maximal prunings relative to an invariant set.

Throughout this section we only consider prunings of the horseshoe. However, much of the theory could equally well be developed for any *Markov thick tree map*: these are maps obtained by

³The two papers [32] and [36] have slightly different definitions of pruning fronts and of the construction of the pruning isotopy. This statement is true for the formulation of [32]. For that of [36], all of the dynamics is destroyed with the possible exception of a single periodic orbit.

‘thickening’ Markov tree maps (for example, the horseshoe is obtained by thickening the full tent map on the interval, and the maps depicted in Figure 12 below are obtained by thickening the tree maps of Figure 13). More details about thick tree maps can be found in [19, 36, 43].

4.1 A full family in dimension 2

As described in Section 2, the horseshoe projects to a self-map of the interval. There is a special type of pruning disk for which the associated pruned maps also project to interval maps. These pruning disks are those which go from the top to the bottom of the horseshoe image as shown in Figure 8: we will refer to them as *vertical* pruning disks. Not every such disk satisfies the pruning conditions: in fact (see Proposition 4.2), the pruning conditions are satisfied precisely if the left hand edge of the vertical disk is at a *kneading coordinate* (i.e. no forward iterate of it lies to its right).

Applying the pruning theorem to this one-parameter family of vertical disks produces a two-dimensional analog of a *full* family of unimodal interval maps (i.e. one which presents all possible kneading invariants). Barge-Martin [6] constructed a family of two-dimensional maps with similar dynamics using the inverse limit of the standard unimodal family.




Figure 8: A vertical pruning disk and its unimodal quotient

If the horizontal coordinate of the C -side (that is, the vertical side) of a vertical pruning disk D is preperiodic, then it is possible to find a value of the parameter a (for which the critical point of the map $x \mapsto a - x^2$ is attracted to the corresponding periodic orbit) so that, if $b > 0$ is chosen sufficiently small, the Hénon map $H_{a,b}$ is conjugate to the pruning of the horseshoe associated to D (this follows from the stability of the one-dimensional map which the Hénon map is a perturbation of when b is small). As the parameter a is varied, however, several authors have noted that the Hénon family goes through bifurcations very different from those in the unimodal family (see, for example, [40, 44, 58, 59]). The *anti-monotonicity* results of Kan-Koçak-Yorke [66] are particularly interesting: they show that, near a homoclinic tangency, the Hénon family goes through infinitely many periodic orbit destroying as well as period orbit creating bifurcations. This contrasts with the monotonicity of periodic orbit creation in the unimodal family [71].

4.2 Quotients of other prunings of the horseshoe

If a pruning disk is not vertical, then the associated pruned map does not project to a self-map of the interval. In this section it will be explained how it still projects to a self-map of a one-dimensional space, namely a tree.

We describe first one of the simplest prunings of the horseshoe which gives rise to a nontrivial tree. The resulting pruned map is called the T-map, since the corresponding tree is a *triod* (that is, a tree with one vertex of valence 3 and three vertices of valence 1). More complicated examples will also be given in this section, but a detailed description of how to obtain the associated tree maps is left to Appendix A.

Example 2 Let F be the horseshoe map. \mathbb{T} and its image under F^2 are shown in Figure 9, together with a disk D and its preimage. The marked point which lies on C is the fixed point of code 1. By considering the fates of forward iterates of C and backwards iterates of E , it is easy to see that D is a pruning disk for F (the argument for E is like that of Example 1 a), while the argument for C is like that of Example 1 b)).




Figure 9: The T-pruning disk

Since D lies inside \mathbb{T} , it inherits a decomposition into vertical segments and a half-disk at its right hand end. In Figure 10, $F^{-1}(D)$ is shown together with the preimage decomposition (the vertical segment through the fixed point has been thickened into a small rectangle to clarify the figure). We *amalgamate* $F^{-1}(D)$ with \mathbb{T} , thereby enlarging the decomposition elements of \mathbb{T} which $F^{-1}(D)$ intersects. The resulting space is shown in Figure 11. Its quotient obtained by collapsing the decomposition elements to points is a tree with a valence three vertex. It should be thought of as being obtained from the interval by gluing together the points which are joined by a curved line in the interval beneath Figure 9.

Up to this point nothing has been done to the horseshoe map. We have merely decided to look at it from a somewhat unorthodox point of view, as a map which projects to a self-map of a triod as shown in Figure 13 (center diagram). The advantage is that, with this new decomposition, $F^{-1}(D)$ projects to a segment K whose image *backtracks* (i.e. its image traverses this edge in one direction




Figure 10: The induced decomposition of $F^{-1}(D)$




Figure 11: Amalgamating \mathbb{T} with $F^{-1}(D)$

and immediately re-traverses it in the other direction; see Figure 13, top of center diagram). As stated above, pruning D away is the same as pruning $F^{-1}(D)$ away, and this is shown in Figure 12. On the quotient tree, this corresponds to pulling tight the image $f(K)$ of the backtracking segment, as shown in Figure 13 (diagram on the right).




Figure 12: Pruning $F^{-1}(D)$ away to produce the T-map

The complete process on the quotient tree is depicted in Figure 13. We first identify points between $0\bar{1}$ and $\bar{1}$ which have the same image under the tent map f . This identification turns the interval into a triod; and those arcs of the interval which were mapped by the tent map over the interval between $0\bar{1}$ and $\bar{1}$ now have images which backtrack over the new edge of the triod. There are two such arcs: one from $00\bar{1}$ to $0\bar{1}$, and the other from $1\bar{1} = \bar{1}$ to $10\bar{1}$. The latter arc (denoted K in Figure 13) has *innermost* image (cf. Figure 12); pulling tight the backtracking segment of the image of K yields the quotient of the pruned map of Figure 12.




Figure 13: The T-pruning as seen on the quotient trees

On the symbolic level, the dynamics of the T-map F_D is described by the subshift obtained from $\sigma: \Sigma \rightarrow \Sigma$ by removing the orbits of points $\beta \cdot \alpha$ for which $\alpha \succ \bar{1}$ and $011\bar{0} \prec \beta \prec 111\bar{0}$ (since $\bar{1}$ is the symbolic coordinate of the C -side of D , and $011\bar{0}$ and $111\bar{0}$ are the symbolic coordinates of the E -side). With the exception of orbits of points on the E -side of D (which are not destroyed by the pruning), this is the subshift in which the words 1010 and 1110 are prohibited. For let $x \in \Lambda$ have itinerary $k(x) = \beta \cdot \alpha$. Then x lies in the interior of D or on E if and only if $\alpha \succ \bar{1}$ and

$011\bar{0} \preceq \beta \preceq 111\bar{0}$. The condition on α says precisely that $\alpha = 1^{2l+1}0\dots$ for some $l \geq 0$; and the condition on β is equivalent to $\beta = {}_1^0c$, where $c \succeq 11\bar{0}$, i.e. $c = 1\dots$. Hence the orbit of x intersects $\text{Int}(D) \cup E$ if and only if $k(x)$ contains a word of the form $1_1^0 1^{2l+1}0$; but this is equivalent to its containing either 1010 or 1110.

Notice that this symbolic dynamics is in terms of the original symbolics on the full horseshoe, rather than of a new Markov partition for the T-map. The ability to transfer horseshoe symbolics to pruned horseshoes is an important feature of pruning theory. A more general approach to this procedure is described in Appendix B.

Notational interlude

Pictures of pruning disks and their iterates can become very complicated (as seen in the preceding example, which is one of the simplest prunings yielding a tree other than the interval). It is therefore useful to develop a clearer way of depicting them: this is achieved by eliminating the space between \mathbb{T} and its image. Most of the figures depicting the horseshoe in the remainder of the paper contain a square — the *symbol square* — which represents the nonwandering set of the horseshoe after all of the gaps in the horizontal and vertical Cantor sets have been collapsed. Points in the symbol square are specified by elements of $\{0, 1\}^{\mathbb{Z}}$ (Figure 2). This leads to ambiguity for some points, namely those whose vertical or horizontal coordinates lie on the boundary of one of the gaps. Explicitly, on the square we cannot distinguish between pairs of points of the form $y01\bar{0}$ and $y11\bar{0}$, or between pairs of the form $\bar{0}10x$ and $\bar{0}11x$ (here, and whenever this notation appears, x will always denote a sequence of 0s and 1s which is infinite to the right, and y will always denote a sequence infinite to the left). We incorporate this ambiguity in our notation, denoting the first pair of points by $y_1^0 1\bar{0}$ and the second by $\bar{0}1_1^0 x$. This ambiguity, which is merely notational/pictorial at this point, turns out to have dynamical significance (see Examples 6 and 7).

In figures containing the symbol square, pruning disks will be depicted with shaded boxes. Since a pruning disk in the horseshoe is determined by the two endpoints of its C -side (or, equivalently, of its E -side), a box representing a pruning disk is determined by two of its corners, namely, those in the interior of the square (which are at the intersection of the stable and unstable segments forming its boundary). We will refer to these as the *inner* corners of the pruning disk to distinguish them from the corners lying on the boundary of the square. We will usually draw pruning disks — and the boxes representing them — extending from some interior point to the right edge of the square and, in this case, the corners determining the disks will lie on the same vertical line. Notice that the preimage of a disk extending to the right of the square and symmetric about the horizontal centre line is a disk extending to the top of the square and symmetric about the vertical centre line (as in Figure 9).

The union of the pruning disks is the pruning front, and the (bi-infinite) orbit of the pruning front under the dynamics is the region where the isotopy of the pruning theorem takes place. The schematic representations of the two pruning disks of Figure 7, and of the pruning disk of Figure 9, are given in Figure 14. The symbol square in these figures is also divided into equally-sized regions (as in Figure 2) to help clarify the exact positions of the boundaries of pruning disks.

Understanding this schematic representation of pruning disks is an essential part of understanding the remainder of the paper, and we encourage the reader to become familiar with it before proceeding. Figure 30 a) and c) show some images and preimages of the pruning disks of Figure 14 a) and c): being able to work out the orbit of a disk in this way is also important. The

program [47] can be used to draw and iterate pruning disks.




Figure 14: A schematic representation of the pruning disks of Figure 7 and of Figure 9

A general pruning disk can be specified by the horizontal and vertical coordinates of its edges. The simplest type of pruning disk, the *vertical pruning disk* discussed in Section 4.1, extends all the way from the bottom to the top of the square, and from some point up to the right of the square. It is thus completely determined by the horizontal coordinate of its left edge, which will be referred to as its *horizontal coordinate*.

Following the treatment of Example 2, we now briefly describe the general method for obtaining the quotient tree associated with a given pruning front: see [36] for a more complete description. Let D be a pruning disk for the horseshoe map F . Suppose that by some amalgamation similar to that of Example 2, we arrive at a new region \mathbb{T}_{New} (decomposed into arcs and disks) with $F(\mathbb{T}_{\text{New}}) \subseteq \mathbb{T}_{\text{New}}$, such that $F: \mathbb{T}_{\text{New}} \rightarrow \mathbb{T}_{\text{New}}$ projects to a tree map $f: T \rightarrow T$ under the projection $p_{\text{New}}: \mathbb{T}_{\text{New}} \rightarrow T$ (which collapses each decomposition element to a point): that is, $p_{\text{New}} \circ F = f \circ p_{\text{New}}$. The crucial observation is that if $K \subset T$ is a segment whose image under f backtracks (such as the arc from $\bar{1}$ to $10\bar{1}$ in Example 2), then pulling $f(K)$ tight corresponds to pruning away all orbits of F which fall into $p_{\text{New}}^{-1}(K)$. Now for the original projection $p: \mathbb{T} \rightarrow I$, preimages $p^{-1}(K)$ of segments $K \subset I$ are vertical strips going from top to bottom of \mathbb{T} or, equivalently, vertical strips going from top to bottom of the symbol square. The amalgamation process glues together points on the quotient tree, and we should therefore expect $p_{\text{New}}^{-1}(K)$ to consist of a finite union of vertical strips from top to bottom of the symbol square.

Turning this reasoning around, in order to *construct* the new quotient tree and tree map we look for $n \geq 0$ such that $F^{-n}(D)$ is represented by a union of vertical strips from top to bottom of the symbol square. The existence of such an n is equivalent to the E -side of D lying on $F^n(\partial\mathbb{T})$. A forward F -invariant region can then be constructed by defining

$$\mathbb{T}_{\text{New}} = \mathbb{T} \cup F^{-1}(D) \cup F^{-2}(D) \cup \dots \cup F^{-n+1}(D),$$

and amalgamating the decompositions of $F^{-j}(D)$ (induced by that of $D \subset \mathbb{T}$) with the original decomposition on \mathbb{T} . The quotient space obtained by collapsing the decomposition elements of \mathbb{T}_{New} to points is T and, since the construction is dynamically invariant (which is why we needed to take the union above to form \mathbb{T}_{New}), $F: \mathbb{T}_{\text{New}} \rightarrow \mathbb{T}_{\text{New}}$ projects to a map $f: T \rightarrow T$. Now the n^{th}

preimage $F^{-n}(D)$ of the pruning disk projects to a piece K of T whose image under f backtracks: pruning D (or equivalently $F^{-n}(D)$) away corresponds to pulling K tight.

In fact, if \mathbb{T}_{New} is not simply connected then some additional steps must be carried out: see Appendix A or [36] for a fuller account.

The next two examples are explained in detail in Appendix A (Examples 13 and 14 respectively).

Example 3 Consider the pruning disk D depicted in Figure 15, which has inner corners $\overline{01_1^0 10} \cdot \overline{10}$ and $\overline{01_1^0 11} \cdot \overline{10}$: the periodic point $\overline{10}$ is therefore contained in the C -side of D , and the orbit of this point is also shown in Figure 15. The E -side of D lies on $F^3(\partial\mathbb{T})$, so that the first two preimages of D (also depicted in the figure) are used in the construction of the quotient tree. Amalgamating $F^{-1}(D)$ with \mathbb{T} corresponds to identifying points in the quotient interval between $0\overline{10} = \overline{01}$ and $1\overline{10}$ which have the same image. Once this identification has been carried out, the image of the interval between $11\overline{10}$ and $\overline{10}$ also backtracks, and identifying points in this interval with the same image corresponds to amalgamating $F^{-2}(D)$ with $\mathbb{T} \cup F^{-1}(D)$. The resulting tree map (after pulling tight) is shown in Figure 16. This pruned map is called the *TT-map*.




Figure 15: The TT-pruning disk



Figure 16: The quotient of the TT-map

Example 4 In this example, the pruning front is made up of two pruning disks. The first, shaded lighter in Figure 17, is just the pruning disk of Example 2, with inner corners $\overline{01_1^0 0} \cdot \overline{1}$ and $\overline{01_1^0 1} \cdot \overline{1}$. The second has inner corners $\overline{01_1^0 0100} \cdot \overline{100}$ and $\overline{01_1^0 0101} \cdot \overline{100}$. The quotient tree and tree map are shown in Figure 18. The fixed point of code 1 (marked with \bullet), and the period three orbit of code 100 (marked with \blacksquare) project to the fixed and period 3 valence 3 vertices of the tree.




Figure 17: A pruning front made up of two pruning disks



Figure 18: The quotient tree map from the pruning front of Figure 17

Remark: Not every pruning disk has its E -side on some image of $\partial\mathbb{T}$ (i.e. on the unstable manifold of the fixed point of code 0). In such cases the quotient tree may have infinitely many edges. For example, the quotient tree map associated to Example 1 b) is shown in Figure 19.




Figure 19: The quotient tree map from the pruning disk of Figure 14 b)

4.3 An example in the Hénon family

Here we discuss an example in the Hénon family for which it is possible to compute what seems to be the associated pruning front (so that the dynamics of the Hénon map consists precisely of the horseshoe orbits which don't fall into the pruning front). Davis, MacKay, and Sannami [30] studied the area-preserving Hénon map with parameters $a = 5.4$ and $b = 1$, and suggested that it appears to be hyperbolic (notice that the sign of the parameter b in their formula for the Hénon map differs from the one used here). They supported this claim by proposing a Markov partition

for the map, and showing that the numbers of periodic orbits of each period up to 20 predicted by this partition are identical with those found using the method of Biham and Wenzel [20]. Here we use the pruning front for this example (which itself is given by the analysis in [30] of the excluded orbits in $H_{5,4,1}$) to construct an associated tree map, which in turn yields a Markov partition for the pruned horseshoe. Although this partition is different from that of [30], it can be shown to generate the same subshift of finite type.

Figure 20 shows the stable and unstable manifolds of the outer fixed point of $H_{5,4,1}$. Either by comparing this figure with the configuration of stable and unstable manifolds in the horseshoe (Figure 3 c)), or by reference to [30], we arrive at a (conjectural) associated pruning front, consisting of a single pruning disk D , as shown (with 2 images and 4 preimages) in Figure 21. The inner corners of D have coordinates $\bar{0}1_1^0 010 \cdot 10_1^0 \bar{1}0$ and $\bar{0}1_1^0 011 \cdot 10_1^0 \bar{1}0$.




Figure 20: Stable and unstable manifolds of the Hénon map $H_{5,4,1}$




Figure 21: Conjectural pruning front for the Hénon map of Figure 20

The associated tree map is shown in Figure 22: it has transition matrix

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix}.$$

The symbolic representations — in terms of the original 0s and 1s — of orbits of this map are given by bi-infinite walks on the graph of Figure 23. In Appendix B it is explained how this Markov graph can be obtained from the tree map of Figure 22.




Figure 22: The quotient tree map from the pruning front of Figure 21




Figure 23: The Markov graph giving the symbolic representations of the orbits of $H_{5.4,1}$

4.4 Efficient tree maps

This section introduces some of the tools which will be needed in Section 4.6 to determine maximal prunings relative to periodic and homoclinic orbits of the horseshoe. Given such an orbit P , the problem is to find a maximal pruning front which is disjoint from P (or, in other words, to destroy

as much dynamics as possible while leaving P untouched). Only in the simplest cases can such maximal pruning fronts be found (and be shown to be maximal) by explicit calculation of disks in the symbol square. Efficient tree maps are a useful tool in a more general approach.

Throughout this section, let $T \subset \mathbb{R}^2$ be a tree and $f: T \rightarrow T$ be a self-map which can be ‘performed’ in the plane: that is, f can be obtained by deforming T inside the plane and placing it back onto itself. The tree T is allowed to have valence 2 vertices, and we require that f map vertices to vertices, that it have finitely many critical points (at which f is not locally one to one) and that these critical points also map to vertices. All of the examples discussed in this paper, except for that of Figure 19, are of this type.

Bestvina and Handel made use of the concept of *efficient* tree maps in [19], where they presented an algorithmic proof of Thurston’s classification theorem for surface homeomorphisms up to isotopy (this theorem will be discussed in Section 6.1). Here we give a definition of efficiency adapted to the setting described in the previous paragraph. Let $f: T \rightarrow T$ be a tree map, and P be a periodic orbit of f . Then f is said to be *efficient relative to P* if i) each end of T is a point of P ; and ii) the images of the critical points of f ‘wrap around’ points of P in such a way that it is impossible to shorten the images of the edges of T under iterates of f without crossing over some point in P .

Example 5 In Figure 24 three tree maps are shown, with a periodic orbit P marked on each with small circles (in fact all of the trees are intervals). Some iterates of each tree map are also shown (with the edge images slightly separated, so that it is possible to see what is happening). The first two are efficient relative to P (a method for showing this will be described next), but the third is not: for the third iterate of the arc between successive points of P which contains the point marked x can be pulled tight without passing through any point of P .




Figure 24: Iterates of three tree maps

Efficiency can be detected [49] by tracing back from the critical points of f to see whether or not there is a point of T which is not in the orbit P , and which will become ‘innermost’ in a fold after some number of iterations. (Here the word innermost refers to the way that the edges are mapped as we perform f on the plane.) For example, consider the map of Figure 24 b). In Figure 25 we mark with 0 the top of point c : for c is a critical point, and any arc which passed above c would be mapped by f to an arc which is innermost in the fold (the image of such an arc is also depicted in the figure). In $f(T)$ there are no such arcs. However, an arc passing below d

would be mapped by f to an arc passing above c , and we therefore mark the bottom of d with -1 : any arc passing below d would be mapped by f^2 to an arc innermost in the fold. Again, there are no such arcs in $f(T)$. Finally, an arc passing below b would be mapped by f to an arc passing below d , so we mark the bottom of b with -2 : any arc passing below b would be mapped by f^3 to an arc innermost in the fold. Once again, there are no such arcs in $f(T)$: moreover, since the preimage of b is a , which is an end of T , no such arc can arise in any iterate of f . We therefore conclude that f is efficient relative to P .




Figure 25: Checking that the map of Figure 24 b) is efficient

Now let us carry out the same process for the example of Figure 24 c) (see Figure 26). We mark the top of c and the bottom of e with 0 and -1 respectively. However, there is an arc in $f(T)$ which passes below e , and hence the second iterate of this arc (i.e. the third iterate of an arc of T) is innermost in the fold. Thus f is not efficient relative to P .

Intuitively, this means that f can be deformed so as to make the image under f^3 of the edge containing x shorter. It will be explained in the next section how this can be achieved by pruning away a pruning disk in a two-dimensional map corresponding to f .




Figure 26: Checking that the map of Figure 24 c) is not efficient

4.5 Pruning and efficiency

As stated at the end of the last subsection, there is a relationship between pruning and (the lack of) efficiency. This relationship will now be made explicit. Let $f: T \rightarrow T$ be a tree map that can be performed in the plane. There is an obvious way of constructing a plane homeomorphism which has f as a quotient: simply ‘thicken’ the tree T to a ‘thick tree’ \mathbb{T} and construct a ‘thick tree map’ $F: \mathbb{T} \rightarrow \mathbb{T}$ so that it projects to f (the horseshoe map and the thick ‘Y’ of Figure 12 are examples

of thick tree maps). Figure 27 depicts a thickening of the tree map of Figure 24 c): there are ‘junctions’ above each point of P , which get contracted by F , and the complementary rectangles get expanded horizontally and contracted vertically in standard Markov partition style. Because of this construction, periodic orbits of f and F are in 1-1 correspondence and (as in the case of the horseshoe) we do not distinguish between them. The crucial observation is that f is not efficient relative to P if and only if there exists a (nontrivial) pruning disk D for F which does not contain points of P , so that some of the dynamics of F can be destroyed while leaving P untouched.




Figure 27: A tree map and its thickening

Let us return to the horseshoe and its periodic orbits or, equivalently, to the full tent map $f: I \rightarrow I$ and its periodic orbits. If P is any periodic orbit of f other than the fixed point of code 0 then, since the endpoints of I are not points of P , f cannot by definition be efficient relative to P . This, however, is a minor issue, since there is an obvious way of remedying the situation: simply consider the unimodal map with critical value equal to the rightmost point of P and then cut off the ends to the right of the rightmost point of P and to the left of the leftmost point of P . This new map is denoted by f_P . In the language of one-dimensional dynamics the remaining interval is the *dynamical interval* of the unimodal map whose critical point belongs to P . The two-dimensional thick tree map associated to f_P is dynamically equivalent to the map obtained from the horseshoe by pruning the vertical pruning disks to the right of the rightmost point of P and to the left of the leftmost point of P , as shown in Figure 28. These types of prunings to adjust the ends of a (thick) tree will be done, sometimes without comment, at the start of every example: we refer to them as the *obvious vertical prunings*. After these obvious vertical prunings, we obtain a quotient tree map which satisfies condition i) in the definition of efficiency relative to P .

Figure 29 shows the obvious vertical prunings relative to the horseshoe periodic orbits of codes 10010 and 100110. In fact, only the vertical prunings on the right of the square are shown: the iterates of the right vertical pruning disk contain all nonwandering points in the left vertical pruning disk (except for the attracting fixed point lying outside of S), rendering the left disk essentially irrelevant. We therefore always omit the left obvious vertical pruning disk. The quotient tree maps (before carrying out the additional pruning for 100110 which is also indicated in the figure) are




Figure 28: A period 6 horseshoe orbit, and the two obvious vertical pruning disks relative to it those of Figure 24 b) and c) respectively.




Figure 29: Obvious vertical prunings for the horseshoe orbits 10010 and 100110, and an additional pruning for 100110

The tree map corresponding to 10010 is efficient. It is not hard to convince oneself — and this calculation will be carried out in Example 8 — that any pruning disk for the associated two-dimensional map which does not contain points in the orbit 10010 cannot contain any nonwandering points. The tree map corresponding to 100110, on the other hand, is not efficient. This is reflected by the fact that there is a pruning disk for the corresponding two-dimensional map (also shown in Figure 29) which contains nonwandering points, but is disjoint from the orbit 100110.

This motivates the concept of *maximal* pruning relative to an invariant set. Although the invariant sets we have been talking about thus far have always been periodic orbits, we will be interested in considering other kinds of invariant sets such as collections of periodic orbits and

homoclinic orbits. Let $F: \mathbb{T} \rightarrow \mathbb{T}$ denote the horseshoe (or any other thick tree map) and let Λ be an F -invariant set (i.e. $\Lambda = F(\Lambda) = F^{-1}(\Lambda)$). We say that a pruning front \mathcal{F} (made up of a union of pruning disks) is *maximal* relative to Λ if the pruned map $F_{\mathcal{F}}$ (obtained from F using the more general version of the pruning theorem for a pruning front made of a union of several disks) does not have any nontrivial pruning disks which do not intersect Λ . In other words, \mathcal{F} is maximal relative to Λ if, after pruning \mathcal{F} away, no other pruning is possible unless we also prune part of Λ away.

For the periodic orbit 10010 the obvious vertical prunings were maximal, but this is not the case for the periodic orbit 100110. We discuss one more interesting example before considering the issue of maximal prunings more generally in the next subsection.

Notice that (some iterate of) every nontrivial pruning disk for the horseshoe must contain the centre point on the right edge of the square. This point represents two nonwandering points of the horseshoe whose symbolic representations are $\overline{0}1_1^0 \cdot 1\overline{0}$. It follows that the maximal pruning relative to either of these orbits is empty. In other words, if we agree not to destroy the orbits $\overline{0}1_1^0 \cdot 1\overline{0}$, then we cannot prune anything away from the horseshoe. This is closely related to a result of Handel [51] (see also [75]), stating that either one of these homoclinic orbits *forces* the full horseshoe: that is, no isotopy relative to the given homoclinic orbit can destroy any of the dynamics of the horseshoe.

4.6 Finding maximal prunings

In this section we explain how to find maximal prunings relative to some simple periodic and homoclinic orbits. Some of the pictures which arise require some effort to understand: the reader is again encouraged to make use of the program [47] to experiment with these and similar examples.

For finite invariant sets of thick tree maps there is an algorithm for finding maximal prunings (this statement is also true for horseshoe orbits which are homoclinic to the fixed point of code 0):

Theorem 4.1 *If F is a thick tree map and P is a finite invariant set, then there is an algorithm for finding a maximal pruning of F relative to P .*

The idea of the algorithm is very simple: repeatedly find and prune away pruning disks which avoid the invariant set P . The proof that this process terminates consists of translating the Bestvina-Handel algorithm for finding invariant train tracks for surface homeomorphisms [19] into the language of pruning. Details can be found in [36].

Let D be a disk in the horseshoe — not necessarily a pruning disk — bounded by a segment of stable manifold and a segment of unstable manifold. As has already been explained, a suitable iterate of D extends to the right hand edge of the square, and is symmetrically placed about the horizontal centre line of the square (its preimage then extends to the top edge of the square, and is symmetrically placed about the vertical centre line). Such a disk D is determined by its inner corners, which, by symmetry, have coordinates of the form $y0 \cdot x$ and $y1 \cdot x$, for some semi-infinite sequences x and y . (The notation reflects the fact that the future of a doubly-infinite sequence determines its horizontal position, and its past determines its vertical position.) We will denote such a disk by $D(y * x)$ (the symbol 0_1 is reserved in this section for the ambiguity of coding of points in the square which lie at the ends of gaps in the horizontal or vertical Cantor sets).

In order to discuss an algorithmic approach to maximal prunings, we need a way of searching for pruning disks. Suppose that we fix y and ask for conditions on x such that $D(y * x)$ is a pruning

disk. Since the C - and E -sides of $D(y * x)$ are segments of stable and unstable manifolds, the only issue is whether or not the disk satisfies pruning condition ii), that for all $n > 0$

$$F^n(C) \cap \text{int}(D) = \emptyset = F^{-n}(E) \cap \text{int}(D).$$

Let us consider some simple examples. First, take $y = \bar{0}$. Then the corners of $D = D(y * x)$ will be of the form $\bar{0} \cdot x$ and $\bar{0}1 \cdot x$. These points lie on the bottom and top of the symbol square respectively, so that D is a vertical disk with horizontal coordinate x . Pruning condition ii) requires that $F^n(C) \cap \text{int}(D) = \emptyset$ for $n > 0$. Since vertical segments map into vertical segments, this is equivalent to requiring that the horizontal coordinate of $F^n(C)$ never lie to the right of x : or, in other words, that all shifts of x be less than or equal to x in the unimodal order. This precisely characterizes kneading sequences of unimodal maps.

Proposition 4.2 *A vertical disk $D(\bar{0} * x)$ in the horseshoe is a pruning disk if and only if x is a kneading sequence.*

The argument above takes care of the ‘only if’ statement. The ‘if’ is just as easy.

For a slightly more complicated example, take $y = \bar{0}11$. Figure 30 shows the disks $D(\bar{0}11 * x)$, together with some forward and backward iterates, for three choices of x . Case a) is that of Figure 7 a), together with the iterates $F^{-2}(D)$ through $F^3(D)$. It is clear that, for the iterates shown, pruning condition ii) is satisfied; because further forward and backward iterates of the C - and E -sides, respectively, are contained on the boundary of the square, it follows that the pruning condition is satisfied in general. Case c) is the T-pruning disk $D(\bar{0}11 * \bar{1})$ which we have already encountered. In case b), we can see that $F(C) \cap \text{int}(F^{-1}(D)) \neq \emptyset$ (and hence that $F^2(C) \cap \text{int}(D) \neq \emptyset$), so that this is not a pruning disk. However, it can be enlarged (cf. Figure 6) to give the pruning disk of case c). On the other hand, any disk wider than that of c) (but with $x > 01\bar{0}$) will have the property that $F(C)$ intersects the interiors of both D and its complement (cf. Figure 5). These examples make the general point that, having chosen y , there is a maximal width which can be attained by a pruning disk $D(y * x)$. This maximal width is attained in case c) of the current example: any wider disk must necessarily violate pruning condition ii).




Figure 30: The disk $D(\bar{0}11 * x)$ and some iterates for three choices of x

Let us now consider how to find a maximal pruning relative to an invariant set P . We start by considering some examples where P is a homoclinic orbit. The procedure will always be the same: first choose y to give a disk which is as tall as possible without containing points of P ; then

choose x as far to the left as possible so that $D(y * x)$ is a pruning disk — that is, satisfies pruning condition ii) — and does not contain points of P . Having found one disk, we repeat the procedure, making sure the new disk does not violate the pruning condition with respect to either itself or the previously constructed disks. This turns out to be a rather delicate point when several disks are involved. Interactions between disks lead, in particular, to asymmetric prunings, which will be discussed further at the end of this section.

Example 6 Let P be the homoclinic orbit $\bar{0}1101\bar{0}$ (see Figure 31). We follow the procedure outlined above. Because the point $\bar{0}110 \cdot 1\bar{0}$ on the right hand edge of the square lies in P , in order to obtain a pruning disk which is as tall as possible, we choose $y = y_1 = \bar{0}11$: then for any choice of x , $D = D(y_1 * x)$ touches, but does not contain, the point $\bar{0}110 \cdot 1\bar{0}$. As argued above, taking $x = x_1 = \bar{1}$ makes D into a pruning disk, and any wider disk must necessarily violate pruning condition ii). Therefore $D_1 = D(\bar{0}11 * \bar{1})$ is the first pruning disk in a maximal pruning front relative to P . In fact, for this example, it is the entire maximal pruning front relative to P . Any second disk $D_2 = D(y_2 * x_2)$ would have to be less tall, so as not to contain the point of P on the right hand edge of the square. However, because the preimage of an inner corner of D_1 lies at the centre of the C -side of D_1 (namely, at the point $\bar{0}1_1^0 \cdot \bar{1}$), any choice of x_2 to the left of $\bar{1}$ would mean that D_2 would contain that corner of $F^{-1}(D_1)$ and would therefore violate pruning condition ii). Thus no further pruning is possible and D_1 is a maximal pruning relative to P . Had we chosen P to be any of the four homoclinic orbits $\bar{0}1_{11}^{00}1\bar{0}$, we would have obtained the same pruning disk, except for the usual — and dynamically irrelevant — ambiguity involving the gaps in the Cantor sets. The reason for this is that all four orbits have the same *braid type* (see Section 6).




Figure 31: The maximal pruning front relative to the homoclinic orbit $\bar{0}1101\bar{0}$

Example 7 Now let P be the homoclinic orbit $\bar{0}1001001\bar{0}$ (marked with ■ in Figure 32). Again, because the point $\bar{0}100100 \cdot 1\bar{0}$ of P lies on the right hand side of the square, we choose $y_1 = \bar{0}10010$. The first five preimages of the disk $D_1 = D(y_1 * x)$ are the ones that need to be considered to make sure pruning condition ii) is not violated: for if $n \geq 6$ the E -side of $F^{-n}(D_1)$ is contained in the bottom edge of the square and therefore presents no more danger. The third preimage $F^{-3}(D_1)$ is the closest to the right hand edge of the square and, in fact, it is the one that leads us to choose $x_1 = \bar{1}00$: it is clear from Figure 32 a) that any wider disk would violate pruning condition ii).

Thus $D_1 = D(\overline{010010} * \overline{100})$. Notice that $F^{-3}(D_1)$ has a common boundary segment with D_1 (containing a point of the period 3 orbit of code 100 marked by \bullet in the figure), but the preimages of the inner corner lie quite far from the centre of the C -side of D_1 : this leaves room to add a second disk to the pruning front. We therefore choose y_2 so that y_20 is the vertical coordinate of the third preimage of the inner corners of D_1 , that is $y_2 = \overline{010}$, and start again: look for x_2 so that $D_2 = D(y_2 * x_2)$ does not violate pruning condition ii). At this point we are in the same situation as Example 6, and it is not hard to check that taking $x_2 = \overline{1}$ gives the widest disk which does not violate the pruning condition (see Figure 32 b)). (In fact, this pruning front is exactly that of Example 4.) Thus $D_2 = D(\overline{010} * \overline{1})$. As in Example 6, any of the four homoclinic orbits $\overline{01}_1^0 01_0^0 \overline{10}$ would produce the same pruning front, up to the usual ambiguity. The reason is the same: these orbits have the same braid type.




Figure 32: Pruning disks D_1 and D_2 of Example 7

In summary, the procedure is as follows: choose y_1 to make $D(y_1 * x)$ as tall as possible without containing points of P . Then choose x_1 so that $D_1 = D(y_1 * x_1)$ is as wide as possible without violating the pruning condition and without containing points of P . (In both of the examples above we chose x_1 so as to avoid violating the pruning condition, but we will see an example below in which x_1 is chosen so as to prevent the disk containing a point of P .) Then look for y_2 so $D(y_2 * x)$ is less tall than D_1 and try to find x_2 to the left of x_1 so that $D_2 = D(y_2 * x_2)$ neither violates the pruning condition nor contains points of P . (At this stage, we should be more explicit about what it means not to violate the pruning condition. After the first pruning, we do not have a full horseshoe any longer. In particular, stable and unstable manifolds have changed and it is certainly not obvious that we can choose disks as before, using vertical and horizontal segments inside the square. In Example 7 above and in many other simple examples this is not an issue, but in more complicated examples, a considerable amount of care is necessary.) Repeat this procedure and hope for the best: namely, that after finitely many steps the process stops and it is clear there is nothing else to be pruned away. To be certain that the process has terminated, it is necessary to use the correspondence between pruned maps and tree maps and check that the corresponding tree map is efficient (see Remark b) below).

Let us now consider some examples of maximal prunings relative to periodic orbits.

Example 8 Consider the periodic orbit P with code 10010. As with every periodic orbit, we can perform the obvious vertical prunings. In Figure 33 a), we show the vertical pruning disk $D_0 = D(\bar{0} * \overline{10010})$ and several forward iterates. Notice that the union of the forward iterates contains all points below a certain level y_b , and hence above another level y_t , and therefore also contains a strip in the vertical centre of the square⁴. But the vertical centre of the square is where we should start looking again for further prunings, with y_1 being the vertical coordinate of the point of $F^5(C_0)$ closest to the centre, so as not to violate the pruning condition. Although it is possible to find pruning disks D_1 for this value of y_1 (the maximal such disk is shown in Figure 33 b)), they will not contain any nontrivial dynamics since they are already contained in the central strip which is itself contained in the iterates of D_0 . Therefore, the obvious vertical pruning relative to 10010 is maximal, as claimed earlier.

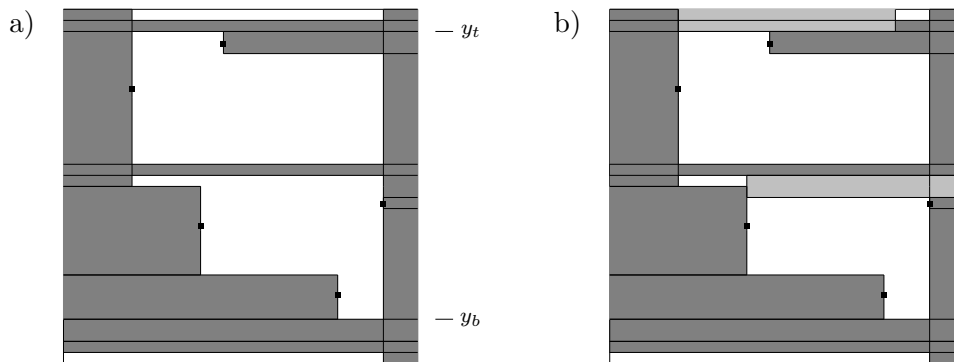


Figure 33: The maximal pruning relative to $\overline{10010}$

Example 9 Next, consider the periodic orbit P with code 1000110. We start with the obvious vertical pruning $D_0 = D(\bar{0} * \overline{1000110})$. The seventh iterate $F^7(C_0)$ of the C -side of D_0 is contained in C_0 , as shown in Figure 34. We choose y_1 so $y_1 0$ is the vertical coordinate of the top of $F^7(C_0)$. A short calculation shows that the widest $D_1 = D(y_1 * x)$ can be while not violating the pruning condition is when $x = x_1 = \bar{1}$. This disk D_1 and its first preimage are shown in Figure 35 a). Notice that, unlike in Example 7, the corner of $F^{-1}(D_1)$ doesn't coincide with the centre of C_1 , the C -side of D_1 . Thus there is room for an additional pruning disk. This time y_2 is chosen so that $y_2 1$ is the vertical coordinate of the inner corners of $F^{-1}(D_1)$, and we seek x_2 so that $D_2 = D(y_2 * x_2)$ is as wide a pruning disk as possible. A rather longer calculation (or the program [47]) gives $x_2 = \overline{0100011}$. (Certainly x_2 couldn't be any further to the left: for if x lies to the left of x_2 then $D(y_2 * x)$ contains the point $\overline{0100011}$ of P ; see Figure 35 b), where D_2 and its preimage are drawn.) The final pruning front made up of three disks D_0 , D_1 , and D_2 is shown in Figure 36.

⁴It is worth mentioning that after the obvious vertical prunings, the 'outer' contour of the set of unstable manifolds of the remaining orbits is made up of the unstable manifolds of a periodic orbit P' . In fact, P' is a periodic orbit of *finite order* braid type (see later definition) and its rotation number is the *height* (again see later definition) of the original orbit P .




Figure 34: The obvious vertical pruning relative to $\overline{1000110}$

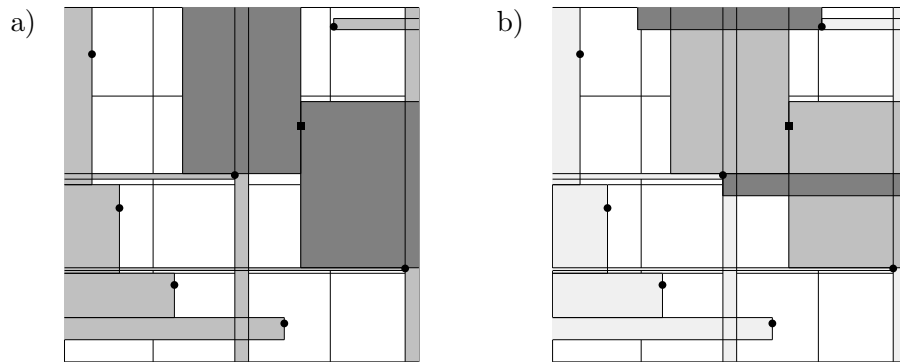


Figure 35: Pruning disks D_1 and D_2 of Example 9




Figure 36: The maximal pruning front relative to $\overline{1000110}$

Remarks:

- a) The general procedure for finding a maximal pruning gave a well-defined result in the simple examples above, but in more complicated cases, when there are many prunings disks, choices will arise. Moreover, there is nothing special about choosing to fix y and then finding the smallest possible x which yields a pruning disk: we could have interchanged the roles of the two coordinates, or applied different strategies at different stages of the algorithm.

In the simplest examples, all of the different possible choices lead to the same maximal pruning front, but this need not always be the case: the maximally pruned maps obtained from different pruning fronts are all conjugate, but are not equal in general.

- b) Except in the simplest cases, it is not easy to decide whether or not the algorithm has terminated (i.e. whether or not any further non-trivial prunings are possible). Not only do the pictures rapidly become extremely difficult to understand, but asymmetric prunings (as described below), which are much harder to analyse, may enter the picture. At present, the only way we are able to prove that the algorithm always terminates (or indeed, to check that it has terminated in a particular case) is by constructing quotient trees and invoking efficiency arguments. It would be interesting to have a purely pruning-theoretic approach to this problem.

When things get interesting: asymmetric prunings

In this section we only describe how asymmetric prunings come about, without going into details of the consequences of this phenomenon, such as the appearance in pruned maps of sinks whose basin boundaries are not rotated. In order to clarify the discussion (which in any case is only on an intuitive level) we ignore the fact that pruned maps are not necessarily differentiable, and assume the existence of stable and unstable manifolds. The important aspects of the discussion do not in fact depend on differentiability.

Figure 4 depicts part of a homoclinic tangle, perhaps of the horseshoe, which contains a pruning disk D . Pruning D away should be thought of as uncrossing the segments of stable and unstable manifolds which lie within D . Figure 37 is a schematic representation of this in the symbol square of the horseshoe.




Figure 37: A schematic representation of Figure 4

Suppose that we have pruned the horseshoe several times, and have arrived at a configuration of pruning disks as shown in Figure 38 a). Figure 38 b) shows the boundary of a disk D bounded

by segments of the stable and unstable manifolds of the pruned map in a). As is clear from the figure, the corners of D are *not* symmetrically placed with respect to the horizontal centre line, even though it may well be the case that, by an appropriate choice of these corners, D satisfies pruning condition ii) and is therefore a pruning disk.




Figure 38: A configuration of pruning disks leading to the possibility of an asymmetric pruning

This situation arises when searching for maximal prunings relative to more complicated orbits than the ones presented above. In the horseshoe, the (pseudo-Anosov) orbits of lowest period for which it occurs are those with codes $10001_1^0 10100101_1^0$. The corresponding pseudo-Anosov maps have a period 2 orbit of 3-pronged singularities with rotation $2/3$, and a period 7 orbit of 3-pronged singularities with rotation 0. (A 3-pronged singularity is a point where three leaves of both the stable and unstable foliations of the pseudo-Anosov map ϕ meet. To say that the period 7 orbit has rotation 0 means that ϕ^7 maps each of the three stable and three unstable leaves to itself.) In fact, fixed-prong singularities can only arise in pruned horseshoes via asymmetric prunings.

5 Families of efficient tree maps

In this section, we motivate and state a conjecture which organizes all horseshoe periodic orbits into families which are linearly ordered by forcing: within families, the forcing order coincides with the unimodal order on the symbolic codes of the orbits. In order to define the families, we describe a recipe for parsing the code c of a periodic orbit into two segments, the *prefix* p and the *decoration* w : these are constructed in such a way that $c = p_1^0 w_1^0$ (the final symbol of c and the symbol between the prefix and decoration are unimportant). The linearly ordered families are precisely the families of orbits with a given decoration: thus the statement that forcing within families coincides with the unimodal order means that $p_1^0 w_1^0$ forces $p_2^0 w_1^0$ if and only if $p_1^0 w_1^0 \succeq p_2^0 w_1^0$. Moreover, the homoclinic orbit $\bar{0}1_1^0 w_1^0 1\bar{0}$ forces all of the periodic orbits with decoration w . These statements can be proved for certain families of decorations [35, 33].

Given a tree map $f: T \rightarrow T$ (once more performed in the plane), it is natural to try to find all of its periodic orbits relative to which a map of the same ‘shape’ as f is efficient. Having the same shape means, for example, being a unimodal (or n -modal) map of the interval: this has obvious generalizations to maps of other trees which are not defined precisely here — what is meant by it in the context of this paper will in any case be clear. In terms of pruning, the question is: what are the periodic orbits P for which the obvious vertical prunings of the corresponding thick tree map $F: \mathbb{T} \rightarrow \mathbb{T}$ form a maximal pruning front with respect to P ? (For a general thick tree map, the ‘obvious vertical prunings’ are prunings to make the ends of the associated tree map into points of the periodic orbit in question. In particular, they need not *look* vertical.)

5.1 The I-efficient orbits

This problem was solved for unimodal interval maps in [48], where all unimodal periodic orbits P for which the associated unimodal map f_P is efficient were found. The solution turns out to be very closely related to the corresponding problems for other tree maps which arise from pruned horseshoes, and we therefore treat it in some detail.

The pruning formulation of the problem is: what are the horseshoe periodic orbits P for which the obvious vertical prunings form a maximal pruning front (we saw in Example 8 that the orbit of code 10010 has this property). We will call these orbits *I-efficient*, since they are efficient on the interval I . They are parameterized by a rational number $q = m/n$ with $0 < q < 1/2$ called *height*, and their codes are of the form $c_{q_1}^0$, where c_q is a particular length $n+1$ word. The simplest way to describe these words is as follows: given $q = m/n$, consider the straight line L in \mathbb{R}^2 from $(0,0)$ to (n,m) . For $0 \leq i \leq n$, let s_i be equal to 1 if L crosses a horizontal line $y = \text{integer}$ for $i-1 < x < i+1$, and equal to 0 otherwise. Then $c_q = s_0 s_1 \dots s_n$. Thus the example of Figure 39 shows that $c_{3/10} = 10011011001$.




Figure 39: $c_{3/10} = 10011011001$

Intuitively, $c_{m/n}$ starts and ends with 1, and between contains m (possibly empty) blocks of 0s, separated by $m-1$ occurrences of 11: the $n-2m+1$ 0s are distributed as ‘even-handedly’ as possible amongst the m blocks. Thus, for example, $c_{1/n} = 10^{n-1}1$, $c_{2/2n+1} = 10^{n-1}110^{n-1}1$, $c_{3/3n+1} = 10^{n-1}110^{n-2}110^{n-1}1$, and $c_{3/3n+2} = 10^{n-1}110^{n-1}110^{n-1}1$.

The words c_q are ordered in the unimodal order inversely to the height q : that is, the periodic point whose itinerary is $\overline{c_{q_1}^0}$ lies to the right of that with itinerary $\overline{c_{q'_1}^0}$ if and only if $q < q'$. The notion of height, both for these particular orbits and for horseshoe periodic orbits in general (see Section 5.2) will be important in much of the remainder of the paper.

In the full horseshoe, (suitable iterates of) pruning disks are always symmetric with respect to the central horizontal line in the square, and it is therefore reasonable to expect that the I-efficient orbits have their rightmost points ‘as close as possible’ to this centre line, in order that it is impossible to do any prunings after the obvious vertical pruning (as in Example 8).

Figure 40 a) is a plot of the rightmost points $c_{q_1}^0$ of the I-efficient orbits: as q decreases, they approach the centre point on the right hand side of the square. As was mentioned in Section 4.5, this point lies on the homoclinic orbit relative to which no pruning of the horseshoe is possible. We shall see that this is a common feature for efficient families. To conform with the notation of [35], we will denote the I-efficient family by \mathcal{D} and will let P_q denote the periodic orbit whose code is (one of) $c_{q_1}^0$. Thus

$$\mathcal{D} = \{P_q : 0 < q < 1/2\}.$$

In addition, we let P_0 denote the homoclinic orbits $\overline{01^0_1 10}$. The observation of the first part of this

paragraph can therefore be restated as follows: the rightmost point of P_q tends to the rightmost point of P_0 as q tends to 0.




Figure 40: The rightmost points of the I-efficient and T-efficient orbits

5.2 Efficiency on other trees

The same problem can be posed for other tree maps arising from prunings of the horseshoe, such as those induced by the T- and TT-maps. Again, this is equivalent to asking: what are the periodic orbits of the associated thick tree maps, with respect to which the obvious vertical prunings are maximal? Since the thick tree map is itself obtained from the full horseshoe via some pruning front \mathcal{F} , another way to state the question is: what are the horseshoe orbits relative to which the maximal pruning front is given by adding the pruning front \mathcal{F} to the obvious vertical prunings?

Figure 41 shows two pruning fronts, each consisting of the T-pruning together with an obvious vertical pruning with respect to a periodic orbit: the periodic orbits in the two pictures are those with codes 1000100 and 100001000, which we denote P_1 and P_2 respectively. It can be checked that P_1 is T-efficient (that is, the left hand picture shows a maximal pruning front with respect to P_1), while P_2 is not (an additional pruning is possible). Intuitively, the reason for this is that the rightmost point of P_1 is very close to the T-pruning disk, thus leaving no room to introduce another (non-trivial) pruning disk; the rightmost point of P_2 is far away from the T-pruning disk.

Figure 40 b) shows the rightmost points of the T-efficient orbits. As for the I-efficient orbits, they tend towards the homoclinic points at the corners of the T-pruning disk on the right hand side of the square, whose symbolic coordinates are $\bar{0}1_{11}^{00} \cdot 1\bar{0}$. The orbits of these points are precisely those relative to which the T-pruning is maximal.

The plot of Figure 40 b) clearly requires us to have calculated the T-efficient orbits: we have also calculated the TT-efficient orbits, and those corresponding to some other tree maps arising from prunings of the horseshoe. Having done this, a striking observation can be made: for each such tree map, the codes of the efficient orbits are of the form $c_q^0 w_1^0$, where the c_q are as defined in Section 5.1, and w is a finite word which depends only on the tree map being considered. In fact, w is precisely the word for which a maximal pruning relative to the homoclinic orbit $\bar{0}1_1^0 w_1^0 1\bar{0}$ gives rise to the tree map being considered. (So, for example, w is the empty word \emptyset for the T-efficient family, and $w = 1$ for the TT-efficient family.) This observation gives rise to an obvious conjecture: before stating it, we fix some notation. Given a periodic orbit P with code c_P , define the *height* $q(P)$ of P to be the rational $m/n \in (0, 1/2)$ with greatest denominator such that c_P starts with




Figure 41: Two pruning fronts made up of the T-pruning together with an obvious vertical pruning

the word $c_{m/n}$ ⁵. Given any word w , let $q_w \in (0, 1/2]$ be the supremum of the values of $q \in (0, 1/2)$ for which all four of the words $c_{q^0}w_1^0$ are codes of periodic orbits of height q . Then if $q < q_w$, let P_q^w denote (any one of) the four periodic orbits of code $c_{q^0}w_1^0$, and set

$$\mathcal{D}^w = \{P_q^w : 0 < q < q_w\}.$$

Moreover, let P_0^w denote (any one of) the four homoclinic orbits whose codes are $\bar{0}1_1^0w_1^0\bar{1}\bar{0}$. (The reader who hasn't considered horseshoe orbits from this point of view may be confused by the use of a single symbol to denote any one of four different periodic (or homoclinic) orbits. The reason for doing so is that the four orbits all have the same *braid type* (see Section 6.1): that is, they are topologically indistinguishable [34].)

The statement of the following conjecture is somewhat problematic, as described in the discussion which follows it: however, it makes a good starting point, since turning it into a conjecture which has a chance of being strictly true involves some linguistic contortions.

(Inaccurate) Conjecture 1 *Let $f: T \rightarrow T$ be the quotient tree map of the maximal pruning of the horseshoe relative to any one of the homoclinic orbits $\bar{0}1_1^0w_1^0\bar{1}\bar{0}$. Then the family of f -efficient periodic orbits is \mathcal{D}^w .*

It is easy to see that, for each fixed w , the rightmost points of the periodic orbits P_q^w converge to the rightmost point of the homoclinic orbit P_0^w as q tends to 0. It can also be shown that the codes $c_{q^0}w_1^0$ are (unimodal) ordered inversely to their height q .

Notice that for every horseshoe periodic orbit P (with the exception of the finite order ones mentioned in footnote 5, and of the I-efficient orbits), there is a word w such that its code can be written in the form $c_P = c_{q^0}w_1^0$, where q is the height of P (this statement is a triviality: by the definition of height c_P starts with c_q , and we just choose w to fit the remainder of the word c_P). Thus the conjecture organises almost all horseshoe periodic orbits into efficient families. We call the word w the *decoration* of P . Thus P is defined (up to the ambiguity involved in the 0 s) by its

⁵Thus we are finding the longest word c_q which is an initial subword of c_P . In fact this definition of height fails for periodic orbits of *finite order* type: this is a small and well understood set of orbits. For motivation, a more precise definition, and an algorithm, see [48].

height and decoration. The decoration describes which family P belongs to, and the height gives its position within that family.

Let us summarize the discussion so far. Given a homoclinic or periodic orbit of the horseshoe, there is an associated maximal pruning front (a collection of pruning disks). Fix a homoclinic orbit $H = \overline{0}1_1^0 w_1^0 \overline{1}0$. Let \mathcal{F}^w denote the maximal pruning relative to H . It cannot contain a vertical pruning disk, since H has a point, namely $\overline{0}1_1^0 w_1^0 \cdot 1\overline{0}$, which lies on the right hand side of the square. If we now ask for the periodic orbits relative to which the maximal prunings are obtained by just adding to \mathcal{F}^w the (obvious) vertical pruning disks through their rightmost points, then the conjecture states that they are precisely those whose codes are of the form $c_q^0 w_1^0$.

This statement is false as phrased: a counterexample is provided, for instance, by the decoration $w = 010$. The maximal pruning relative to the homoclinic orbit P_0^{010} is given in Example 7, and it can be checked that this pruning destroys all of the periodic orbits in the family \mathcal{D}^{010} . The maximal prunings relative to these periodic orbits can therefore certainly not be obtained by starting from the maximal pruning relative to P_0^{010} . Rather than being rejected, however, the conjecture needs to be reinterpreted.

We have already mentioned the important fact that maximal prunings relative to a periodic or homoclinic orbit are not unique in general. (In particular, we should go back and substitute the indefinite for the definite article before every occurrence of the expression ‘maximal pruning’ above.) This non-uniqueness is very difficult to understand completely and it is one of the main obstacles to proving statements such as the (appropriately interpreted) Conjecture 1. However, it is also an essential ingredient in interpreting such statements. In particular, for the (counter-) example above with $w = 010$, it is possible to construct finite approximations to what seems to be a pruning front \mathcal{F}^{010} with infinitely many pruning disks satisfying the properties stated in the conjecture: that is, it is maximal relative to the homoclinic orbit P_0^{010} , does not contain the orbits in the family \mathcal{D}^{010} , and the maximal prunings relative to periodic orbits in \mathcal{D}^{010} are obtained by adding to \mathcal{F}^{010} the obvious vertical disks.

A more palatable way of resolving the situation for individual counter-examples such as $w = 010$ is to use the notion of braid type, discussed in the following section. The orbits in the family \mathcal{D}^{110} have the same braid types (i.e. are topologically indistinguishable) from corresponding orbits in \mathcal{D}^{010} , and are not contained in \mathcal{F}^{010} . This avoids the need to construct a pruning front with infinitely many disks in this example: however, it seems unlikely that this would be a useful tool in a general proof, since the question of when two horseshoe periodic orbits have the same braid type appears to be quite intractable in general.

6 Braid types, forcing and pruning

In this section we describe the notion of the braid type of a periodic or homoclinic orbit [22, 23]. The use of braid types is essential for a proper understanding of forcing; they will also give rise to a more accessible version of the Pruning Front Conjecture, the *Braid Type Conjecture*.

6.1 Braid types and forcing

When working with families of plane homeomorphisms such as the Hénon family, it is usually not clear how to introduce symbolic dynamics for specific maps. Even harder is the problem of doing so coherently for all maps in the family, as one does, for example, with a unimodal family of interval

maps when one decides to assign symbol 0 to points on the left of the critical point and symbol 1 to those on the right. Braid types provide a good way of relating periodic orbits of different maps in the family when one doesn't have access to common symbolics: they are two-dimensional analogues of the permutation of a periodic orbit of an interval map. If we restrict ourselves to periodic orbits of unimodal maps, the permutation determines the symbolic representation, except possibly for one symbol (since the point of the orbit nearest the critical point can usually lie on either side of it without changing the permutation). With braid types, however, this is not the case, even if we restrict our attention to horseshoe braids. That is, the braid type of a horseshoe periodic orbit does not in general determine its code up to some minor ambiguity. This is related to the fact that it is possible to follow a periodic orbit around a loop in the Hénon parameter space, and end up with a different orbit [55]. It is also related to the non-uniqueness of maximal prunings, as was discussed briefly at the end of the last section.

Recall that an isotopy is a continuously varying family $\{f_t\}$ of homeomorphisms depending on a real parameter. Two homeomorphisms f and g are *isotopic* if there is an isotopy starting at one and ending at the other. An isotopy is *relative* to a set Λ if all of the maps in the isotopy agree on Λ .

Being isotopic relative to a set Λ is an equivalence relation. If $F: \mathbb{R}^2 \rightarrow \mathbb{R}^2$ is a homeomorphism and Λ is an F -invariant set, then the *braid type* $\text{bt}(\Lambda, F)$ of Λ is the isotopy class of F relative to Λ up to topological change of coordinates (i.e. up to conjugation by homeomorphisms).

Here is the motivation for the term 'braid type'. Suppose that Λ is a periodic orbit of F (which is now taken to be a disk homeomorphism for the purposes of this motivation). Change coordinates so that the points of Λ are arranged on the horizontal diameter of the disk. Now let $\{f_t\}$ be an isotopy from the identity to F , and consider the paths traced out by $f_t(\Lambda)$ as t varies. Since Λ is a periodic orbit, the set of start points and the set of end points of these paths coincide, and the paths therefore trace out a braid. The braid type of Λ is the conjugacy class of this braid in the Artin braid group (up to full twists).

Example 10 Figure 42 shows the braids corresponding to the horseshoe periodic orbits with codes 1001010_1^0 and 1001110_1^0 (where the points of the orbits have been arranged on the horizontal diameter in the natural way, i.e. preserving their horizontal ordering in the horseshoe). Although they are different braids the associated braid types are the same, since the braids are conjugated in the braid group by the generator σ_4 . Thus the second braid can be obtained from the first by carrying out a coordinate change which interchanges the positions of the fourth and fifth points in the orbit.

Boyland [22, 23] introduced a partial order on the set of braid types of periodic orbits as follows: we say that a braid type β_1 *forces* a braid type β_2 (and write $\beta_1 \geq \beta_2$) if every homeomorphism which has a periodic orbit of braid type β_1 also has a periodic orbit of braid type β_2 . Thus, denoting by $\text{bt}(F)$ the set of all braid types of periodic orbits of the map F , forcing is described by

$$\beta_1 \geq \beta_2 \iff [\forall F, \beta_1 \in \text{bt}(F) \Rightarrow \beta_2 \in \text{bt}(F)].$$

Thurston's classification theorem for isotopy classes of surface homeomorphisms [76] (see also [41, 19]) provides the main theoretical tool for working with braid types. The most important consequence of the theorem for the purposes of this paper can be stated as follows: if F is a plane




Figure 42: Two horseshoe periodic orbits with the same braid type

homeomorphism and P is a finite invariant set (that is, a periodic orbit or a finite union of periodic orbits), then there is a canonical representative in the isotopy class of F relative to P . If the isotopy class is not *reducible* (that is, if it is not possible to group points together and cut the plane into disks which return to themselves after finitely many iterates; *renormalizable* would also be a descriptive name for such maps), then the canonical representative in the isotopy class of F relative to P is either of *finite order* (that is, some iterate of it is the identity), or *pseudo-Anosov* (a map which, like a torus Anosov, preserves a transverse pair of foliations, contracting one of them and expanding the other). Moreover, the Thurston representatives minimize the amount of dynamics in their isotopy classes. This minimality can be expressed in a variety of ways: for example, they minimize topological entropy; and they also minimize $\text{bt}(F)$. The second statement means that, in principle, the problem of deciding which braid types are forced by a given braid type $\beta = \text{bt}(P, F)$ can be solved by finding the Thurston representative in the isotopy class of F relative to P . For if we call this representative Φ , then β forces β' if and only if $\beta' \in \text{bt}(\Phi)$. In practice, this approach is prohibitively time consuming: although it is fairly easy to enumerate the periodic orbits of Φ using train track techniques, determining whether or not a braid corresponding to such a periodic orbit does or does not represent a given braid type requires solving the conjugacy problem in the braid group: the only known algorithms for doing this are so slow as to be totally impracticable except for orbits of very low period.

We have already mentioned that horseshoe orbits with different symbolic representations can have the same braid type. An important question is: how does the number of distinct period n horseshoe braid types grow as n tends to infinity? It seems reasonable to conjecture that the growth is exponential, but as far as we are aware no progress has been made on this problem.

It is possible to generalize some aspects of Thurston's classification to include infinite orbits with tame behaviour towards the ends [50, 42, 24, 62]. In particular, similar techniques can be applied to study forcing among homoclinic orbit braid types.

6.2 Pruning and forcing

Thurston's canonical representative minimizes the dynamics in the isotopy class of a plane homeomorphism relative to a finite invariant set P . Pruning, on the other hand, destroys dynamics by isotopies, and these isotopies are relative to P provided that the interiors of the pruning disks don't intersect P . It turns out that such pruning isotopies are all that is required to remove inessential

dynamics in the isotopy class: in [36] it is shown that if P is a (finite collection of) periodic orbit(s) of a thick tree map F (such as the horseshoe), then a maximal pruning of F relative to P is the Thurston representative of the isotopy class of F relative to P , up to collapsing some wandering domains (i.e. up to a dynamically irrelevant semi-conjugacy: see Section 7.1). In other words, a maximal pruning of F relative to P is essentially equivalent to the Thurston representative of $\text{bt}(P, F)$: it therefore makes it possible to study the braid types which are forced by $\text{bt}(P, F)$.

Here is a more detailed restatement of Conjecture 1:

Conjecture 2 *Given a decoration w , there exists a maximal pruning \mathcal{F}^w relative to the homoclinic orbit $P_0^w = \overline{01_1^0 w_1^0 10}$ which does not contain any point in the family*

$$\mathcal{D}^w = \{P_q^w = c_{q1}^0 w_1^0 : 0 < q < q_w\}.$$

A maximal pruning relative to P_q^w can be obtained by adding to \mathcal{F}^w the obvious vertical pruning through the rightmost point of P_q^w . As a consequence, the family \mathcal{D}^w is linearly ordered under forcing inversely to the parameter q , that is, P_q^w forces $P_{q'}^w$ if and only if $q < q'$.

The statement about forcing follows because, as we saw above, $\overline{c_{q1}^0 w_1^0}$ lies to the right of $\overline{c_{q'1}^0 w_1^0}$ exactly when $q < q'$. Therefore the maximal vertical pruning up to the rightmost point of P_q^w does not contain any point of $P_{q'}^w$ if $q < q'$. It follows from the discussion preceding the conjecture that P_q^w forces $P_{q'}^w$ if $q < q'$.

Examples 11 a) Let $w = \emptyset$ be the empty word. Figure 43 a) depicts the homoclinic orbit $P_0^\emptyset = \overline{010010}$, together with a maximal pruning front \mathcal{F}^\emptyset . In Figure 43 b), the periodic orbits $P_{1/4}^\emptyset$ and $P_{2/7}^\emptyset$ (with codes 1000100 and 1001100100) are shown (depicted \bullet and \circ respectively), together with a maximal pruning front $\mathcal{F}_{1/4}^\emptyset$ relative to $P_{1/4}^\emptyset$, which is precisely \mathcal{F}^\emptyset together with a vertical pruning disk with horizontal coordinate $\overline{1000100}$. Since $2/7 > 1/4$, the rightmost point of $P_{2/7}^\emptyset$ lies to the left of this vertical pruning disk, and hence no point of $P_{2/7}^\emptyset$ falls into $\mathcal{F}_{1/4}^\emptyset$. It follows that $P_{1/4}^\emptyset$ forces $P_{2/7}^\emptyset$.

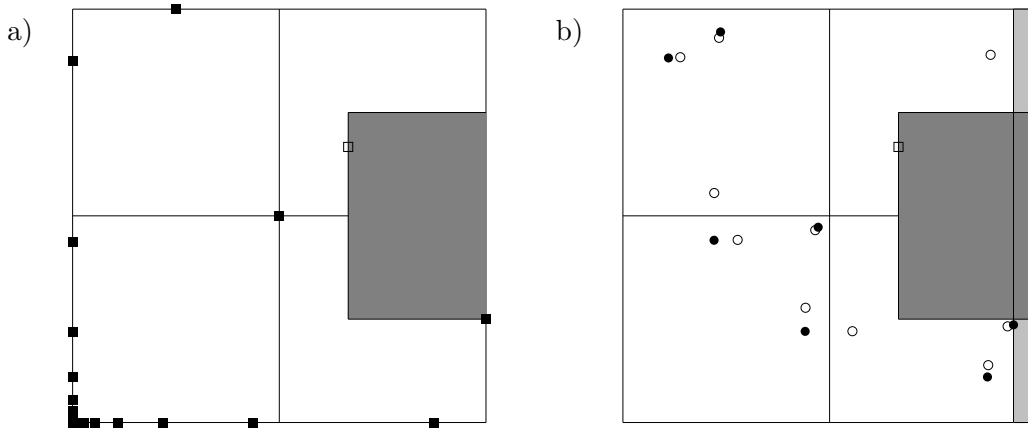


Figure 43: Maximal pruning fronts for the decoration \emptyset

The fixed point of code 1 is also depicted (with \square) in Figure 43 a). The fact that \mathcal{F}^\emptyset has this fixed point on its boundary accounts for the quotient tree corresponding to the decoration \emptyset having a fixed valence 3 vertex, as was seen in Example 2.

- b) A similar treatment of the decoration $w = 110$ yields Figure 44. In a) the homoclinic orbit $P_0^{110} = \overline{01011001\overline{0}}$ is depicted together with a maximal pruning front \mathcal{F}^{110} , consisting of two pruning disks. The fixed point of code 1 and the period 3 orbit of code 110 are also shown: the fact that they lie on the boundary of \mathcal{F}^{110} accounts for the quotient tree corresponding to the decoration 110 having four valence three vertices, one fixed and the others lying on a period 3 orbit (cf. Example 4).

In Figure 44 b), a maximal pruning front $\mathcal{F}_{1/4}^{110}$ relative to the periodic orbit $P_{1/4}^{110}$ with code 1000101100 is shown: it is exactly \mathcal{F}^{110} together with a vertical pruning disk with horizontal coordinate $\overline{1000101100}$.

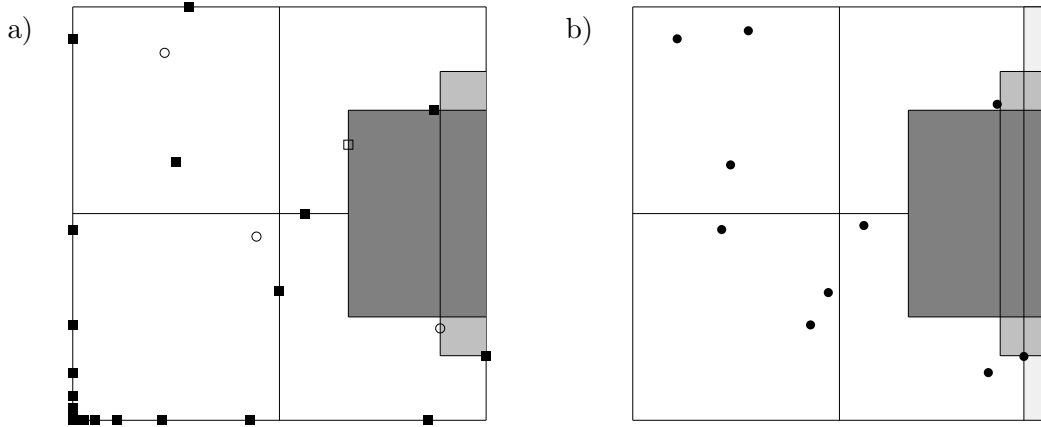


Figure 44: Maximal pruning fronts for the decoration 110

Similar arguments can be used to understand forcing between orbits in different families. Suppose that w and w' are decorations such that $P_0^w = \overline{01_1^0 w_1^0 \overline{0}}$ forces $P_0^{w'} = \overline{01_1^0 w_1'^0 \overline{0}}$. Since for each w , P_0^w forces all orbits in the family \mathcal{D}^w , and since P_0^w forces $P_0^{w'}$ it is reasonable to expect that there exists a maximal pruning front \mathcal{F}^w relative to P_0^w which avoids *both* families \mathcal{D}^w and $\mathcal{D}^{w'}$. Now assume $q < q'$: then $\overline{c_{q_1}^0 w_1^0}$ lies to the right of $\overline{c_{q'_1}^0 w_1'^0}$. Therefore, the maximal pruning front relative to P_q^w — which is obtained by adding to \mathcal{F}^w the obvious vertical pruning for P_q^w by Conjecture 2 — cannot contain points of $P_{q'}^{w'}$ and therefore forces it. This leads to

Conjecture 3 *If $P_0^w = \overline{01_1^0 w_1^0 \overline{0}}$ forces $P_0^{w'} = \overline{01_1^0 w_1'^0 \overline{10}}$ and $q < q'$, then $P_q^w = c_{q_1}^0 w_1^0$ forces $P_{q'}^{w'} = c_{q'_1}^0 w_1'^0$.*

7 The Pruning Front Conjecture

Since Hénon introduced, in 1976, the family of plane diffeomorphisms that now bears his name, mathematicians and physicists alike have tried to understand the prodigiously complicated dynam-

ical behaviours which it presents. The success of Milnor and Thurston’s kneading theory [71] in describing the dynamics of unimodal families of interval maps might lead one to try to develop an analogous theory for families of plane homeomorphisms.

Pruning theory was introduced by Cvitanović as such an analogous theory to describe families — such as those of Hénon and Lozi — which pass from having trivial dynamics to a full horseshoe as parameters are varied. He proposed that, for intermediate values of the parameters, a Hénon or Lozi map is a *partially formed* horseshoe: that is, a horseshoe from which some orbits have been destroyed or *pruned away* (see Figure 45).




Figure 45: Forbidden blocks in a partially formed horseshoe

From the point of view of applications, this approach relies heavily on the assumption that it is possible to introduce symbolic dynamics in a way which is independent of the choice of parameters, in analogy to the fact that in dimension 1 we always use the critical point to introduce symbolic dynamics. Although this is a worthy goal to pursue — indeed, probably the most important goal to pursue if one has applications in mind — the fact that the Hénon family consists of C^ω plane diffeomorphisms, with nothing strongly resembling a critical locus, makes this assumption somewhat shaky ground upon which to try to build a mathematical theory. Notice that the situation is different for the Lozi family where the y -axis, that is, the non-differentiability locus, plays the role of the critical locus. Ishii [63, 64] has used this observation and the piecewise linear nature of the Lozi family to restate and prove the pruning front conjecture in this context.

Here we propose a different approach to this problem, that of giving a complete description of the topological possibilities for the dynamics of a family of two-dimensional homeomorphisms which passes from trivial to rich dynamics as parameters are varied. The idea is to start with the horseshoe and use the Pruning Theorem repeatedly to construct the family of all of its possible prunings. We call this the *pruning family* associated to the horseshoe (of course, one can start with

any other plane homeomorphism and proceed similarly); it is a family of topological models. The problem is then that of comparing the actual systems we are interested in (say the Hénon family) with the models we have constructed. The goal would be to show that, in some reasonable sense, we have constructed enough models to describe every map in the Hénon family. Notice that the pruning family depends on infinitely many parameters: we can prune away a finite but arbitrary number of disks and, even though their C - and E -sides cannot be chosen freely, they are nontrivial parameters that may be varied. Therefore, if we manage to show the pruning family contains the Hénon family, the latter will be only a two-dimensional subspace of an infinite dimensional space.

7.1 The 0-entropy equivalence relation

In this section we describe a way of modifying two-dimensional maps by a semi-conjugacy which collapses ‘irrelevant’ dynamics: that is, parts of the space which do not carry entropy. The semi-conjugacy is defined quite generally and the space of maps it yields is quite interesting in its own right. A more thorough treatment can be found in [31, 37]. It is in this class of maps that it is proposed to compare topological models with the systems that we hope they model.

Recall Bowen’s definition of topological entropy [21]. If X is a metric space with metric d and $F: X \rightarrow X$ is a continuous map, we say that a subset A of X is (n, ϵ) -separated if it is possible to distinguish between the orbits of any two points $x, y \in A$ up to $n - 1$ iterates with precision ϵ : that is, $d(F^j(x), F^j(y)) > \epsilon$ for some $0 \leq j < n$. The topological entropy of F is defined to be the limit as $\epsilon \rightarrow 0$ of the exponential growth rate of the maximum size of an (n, ϵ) -separated subset as $n \rightarrow \infty$. If $K \subset X$ is a fixed compact subset and we only count those orbits which start in K , we obtain the entropy of F in K , denoted $h_F(K)$. Thus

$$h_F(K) = \lim_{\epsilon \rightarrow 0} \limsup_{n \rightarrow \infty} \frac{1}{n} \log (\text{maximal cardinality of an } (n, \epsilon)\text{-separated subset } A \subseteq K)$$

If F is a homeomorphism, we define two points x and y to be *0-entropy equivalent* if there is a compact connected set K which contains both points and for which

$$h_F(K) = 0 = h_{F^{-1}}(K).$$

Remark: If K is a proper subset of X , then it is not necessarily the case that $h_F(K) = h_{F^{-1}}(K)$.

This equivalence is most interesting for two-dimensional systems.

Example 12 Here we describe the 0-entropy equivalence classes for the horseshoe map F . Denote by \mathcal{H}^u and \mathcal{H}^s the closures of the unstable and stable manifolds of the fixed point of code 0 (or indeed of any other periodic point, since their closures coincide) and let $\mathcal{H} = \mathcal{H}^s \cup \mathcal{H}^u$. Equivalence classes are of four kinds:

- a) Closures of connected components of $\mathbb{R}^2 \setminus \mathcal{H}$.
- b) Closures of connected components of $\mathcal{H}^u \setminus \mathcal{H}^s$ (not already contained in sets in a)).
- c) Closures of connected components of $\mathcal{H}^s \setminus \mathcal{H}^u$ (not already contained in sets in a)).

d) Single points which are in none of the sets in a), b) or c).

In [37] it is shown that if F is a $C^{1+\epsilon}$ surface diffeomorphism, then the 0-entropy equivalence classes form an *upper semi-continuous* decomposition of the surface into continua (i.e. compact connected sets). The main idea of the proof is that if a continuum carries positive entropy then it must intersect a Cantor set's worth of stable or unstable manifolds of a periodic orbit which contains transverse homoclinic intersections (this follows from Katok's results and Pesin Theory [67, 68]). Since the ambient space is a surface, it follows that carrying positive entropy is stable under small perturbations of the continuum. Thus, if a sequence of continua carrying 0 entropy converges (in the Hausdorff sense), then the limiting continuum must also carry 0 entropy: this is equivalent to saying that the 0-entropy equivalence classes form an upper semi-continuous decomposition of the surface.

From the upper semi-continuity of the decomposition, it follows immediately from results of Moore [72] and Roberts-Steenrod [73] that the quotient space is a *cactoidal surface* (roughly speaking, a surface with nodes). Since the equivalence is dynamically defined, F projects to a homeomorphism F/\sim on the quotient space and since the equivalence classes do not carry entropy, it follows from a result of Bowen [21] that F and F/\sim have the same entropy. Moreover, any nontrivial compact connected subset of the quotient space carries entropy of either F/\sim or its inverse.

The quotient map by the 0-entropy equivalence relation should be thought of as a 'tight' version of the original map in which all the wandering domains have been collapsed to points. The quotient of the plane by the 0-entropy horseshoe equivalence of Example 12 is shown in Figure 46. The quotient space is a sphere (since everything outside the homoclinic tangle is collapsed to a point), obtained by identifying the solid boundary in the figure along the dotted semi-circular arcs from the mid-point at the top to the corner point on the lower left. The stable and unstable manifolds of the horseshoe project to two transverse foliations with singularities, represented by solid and dashed lines, respectively. These foliations carry transverse invariant measures whose product gives a measure on the sphere. The quotient map preserves both foliations, dividing one of the transverse measures by 2 and multiplying the other also by 2, so that the product measure is invariant. This map is a *generalized pseudo-Anosov map*: it has all the defining characteristics of a pseudo-Anosov map, except that it has infinitely many singularities (and an accumulation point of these singularities).

7.2 More conjectures

We can now state a version of the Pruning Front Conjecture. It addresses one aspect of the comparison between the topological models and the actual maps in the Hénon family.

Pruning Front Conjecture *Let F be the horseshoe and $H_{a,b}$ be the Hénon family with $0 < b < 1$. Then, up to taking quotients by the 0-entropy equivalence relation, the Hénon family is contained in the pruning family of F . More precisely, let $\mathcal{P}(F)$ denote the closure of the family of all possible prunings of F in the C^0 topology and let \sim denote the 0-entropy equivalence relation. Then, $\{H_{a,b}/\sim; 0 < b < 1\}$ is contained in $\{\Phi/\sim; \Phi \in \mathcal{P}(F)\}$ up to topological conjugacy.*

A possibly weaker version is the following Braid Type Conjecture:

Braid Type Conjecture *The braid types that are present in any Hénon map are always horseshoe braid types. More precisely, if P is a finite $H_{a,b}$ -invariant set for some b with $0 < b < 1$, then there*




Figure 46: The quotient of the plane under the 0-entropy equivalence relation for the horseshoe

exists a finite F -invariant set P' for which

$$\text{bt}(P, H_{a,b}) = \text{bt}(P', F).$$

If BTC can be proved, there is a possible approach to proving PFC as follows. Under suitable transitivity assumptions, taking an exhaustion of the set of periodic (and perhaps homoclinic and heteroclinic) orbits of a $C^{1+\epsilon}$ diffeomorphism H and their associated Thurston representatives should give the quotient map H/\sim in the limit. If H is a Hénon map, and all Hénon braids are horseshoe braids, then the same Thurston representatives can be obtained as prunings of the horseshoe. Therefore, up to 0-entropy collapsing, their limit is contained in the pruning family.

If all the pieces in this approach can be made to work, it would extend to a rather complete ‘synthetic’ description of more general 2-dimensional dynamical systems. If it is known that the periodic orbits of a $(C^{1+\epsilon})$ diffeomorphism H belong to some nice set of braid types (such as the set of braid types of some thick tree map F), then a topologically faithful copy of H can be produced from prunings of F . From the construction described above, the data needed to do this may seem to be the set of all periodic orbits of H . It seems likely that in general much less than this is really needed. Recall Handel’s theorem, that if the homoclinic orbit $\bar{0}111\bar{0}$ is present then so must be all of the dynamics of the horseshoe. This homoclinic orbit is precisely the orbit of 1-pronged singularities in Figure 46 which, in turn, comes from the orbit of innermost ‘bigons’ bounded by arcs of stable and unstable manifolds in the horseshoe. In this example the homoclinic orbit plays a role analogous to that of the critical point for an interval map. In other more complicated examples (such as prunings of the horseshoe) one may need several such ‘critical’ orbits: but, again, they are likely to be the orbits of innermost bigons or, equivalently, the orbits of 1-pronged singularities in the 0-entropy quotient map.

8 Conclusions and problems for further research

It has been shown how to make precise sense of pruning and of pruning fronts; how to realize prunings using the pruning theorem; and how pruning corresponds to a deformation theory for tree maps, which might be described as a *generalized kneading theory* for tree and graph endomorphisms. Here the topological aspects of this generalized kneading theory have been discussed. For a definition of kneading determinants for tree maps and their relationship with the ζ -function, see [4, 1]; for a semi-conjugacy to a piecewise linear model, see [5]. The correspondence between the pruned map and a tree map makes it possible to give a complete description of the symbolic representations of the orbits that survived pruning in terms of the symbolic dynamics of the starting system (0s and 1s in the case of the horseshoe).

A method for finding a maximal pruning which does not destroy a given set of (periodic or homoclinic) orbits has been described. This is intimately connected with Thurston's classification of surface homeomorphisms up to isotopy: indeed, up to collapsing some dynamically trivial domains (by the 0-entropy equivalence relation), a maximal pruning of F relative to a finite invariant set P is the Thurston canonical representative in the isotopy class of F relative to P . Making precise sense of this statement uses the concepts of efficient tree maps and braid types.

Braid types and efficient tree maps are central concepts needed to formalize a surprising conjecture which organizes all horseshoe periodic orbits into families linearly ordered by forcing. Whereas in dimension 1 forcing organizes all unimodal orbits linearly, in dimension 2 the situation is more complicated. Periodic orbits are organized into families according to their decoration — a symbolic word at the end of the code of the orbit, the length of which is determined by a simple algorithm. Among orbits with same decoration, forcing is given by the unimodal order. It seems remarkable that, within the complexity of the two-dimensional forcing order, these linearly ordered families should arise from the symbolics in such a simple and natural way. That the ordering within families coincides with the unimodal order is no coincidence: if w is a decoration, the maximal pruning relative to any of the four homoclinic orbits $P_0^w = \overline{0}1_1^0 w_1^0 1\overline{0}$ projects to a 'unimodal' tree map and the family of periodic orbits with decoration w is a special subfamily of the periodic orbits of this tree map. Similar arguments lead to a conjecture about forcing between families: if P_0^w forces $P_0^{w'}$, then an orbit with decoration w forces one with decoration w' if and only if the former is greater than the latter in the unimodal order. These conjectures fall short of a complete description of forcing between horseshoe orbits: in particular, if neither of P_0^w and $P_0^{w'}$ forces the other, then it is possible for an orbit of decoration w to force some, but not all, of the orbits of decoration w' which are smaller than it in the unimodal order. We also have no good understanding of the relationship between decorations and the structure of their associated tree maps.

The existence of asymmetric pruning disks for the horseshoe was mentioned briefly. This is closely related to the existence of attracting periodic orbits for Hénon maps whose basin boundaries have rotation number 0, or, analogously, to the existence of fixed-prong singularities in pseudo-Anosov maps associated to horseshoe braid types. It would be interesting to have a clearer understanding of these phenomena.

The 0-entropy quotient of a surface homeomorphism was introduced: two points are equivalent if there is a compact connected set joining them which carries no entropy under the map or its inverse. This was used to give a precise mathematical statement of the Pruning Front Conjecture: up to taking 0-entropy quotients, any Hénon map is conjugate to a pruning of the horseshoe. A related statement is the Braid Type Conjecture: the braid types of periodic orbits of Hénon maps

are all horseshoe braid types. The latter statement is a consequence of the former, and is potentially easier to prove: it is a topological statement about how the points of a periodic orbit of a Hénon map ‘braid around’ one another as the map is performed. We have given an outline of how a proof of the Braid Type Conjecture might lead to a proof of the Pruning Front Conjecture.

Here again there are many interesting open questions, besides the conjectures themselves. The family of prunings of the horseshoe depends on infinitely many parameters, whereas the Hénon family depends on only two, and one could study how these parameter spaces are related. The Hénon family can also be studied from the complex point of view: periodic orbits disappear from the real plane because they migrate into $\mathbb{C}^2 \setminus \mathbb{R}^2$. This process is not fully understood. Moreover, the theoretical framework that would enable one to derive the Pruning Front Conjecture from the Braid Type Conjecture applies to the study of 2-dimensional dynamics in a broad context. For example, it is closely related to the existence of piecewise linear models for maps in dimension 2 and to the construction of a semi-conjugacy to the piecewise linear model. Such models include Thurston’s pseudo-Anosov maps but are not restricted to them [31].

A Appendix: From pruning fronts to tree maps

This appendix elaborates on the way in which the tree map corresponding to a given pruning of the horseshoe is constructed: Examples 13 and 14 provide detailed treatments of the pruning fronts considered in Examples 3 and 4.

Recall that $D(y * x)$ denotes the pruning disk whose inner corners are the points $y0 \cdot x$ and $y1 \cdot x$, and which extends to the right edge of the symbol square. In the following discussion it is assumed that the horizontal (E -) sides of the pruning disk $D = D(y * x)$ are contained in the unstable manifold of the fixed point of code 0, so that $y = \bar{0}1w$ for some word w of length $n - 1$: thus $F^{-n}(D)$ has its E -side on the boundary of the symbol square (and hence $F^{-n}(D)$ is a union of vertical strips going from top to bottom of the symbol square).

The general procedure, as described in Section 4.2, is to amalgamate in turn the preimages $F^{-1}(D), \dots, F^{-n+1}(D)$ with the initial thick tree \mathbb{T} , which in this case is just the region on which the horseshoe is defined (cf. Figure 1). Each amalgamation changes the thick tree structure of \mathbb{T} , and hence changes the quotient tree T : certain points in T whose images are the same are identified. For example, the initial quotient tree map $f: T \rightarrow T$ is just a full tent map on the interval $T = I$. Amalgamating $F^{-1}(D)$ with \mathbb{T} has the effect of identifying each pair of points of I with itineraries $\cdot 0z$ and $\cdot 1z$ lying between $\cdot 0x$ and $\cdot 1x$ (i.e. with $z \succeq x$ in the unimodal order). This is illustrated in Figure 47. When all of the disks $F^{-i}(D)$ with $1 \leq i < n$ have been amalgamated, the disk $F^{-n}(D)$ corresponds to a segment K of T whose image backtracks: pulling away this backtracking realises the pruning of D .

In the figures accompanying the examples, the tree T is drawn beneath the symbol square. Initially points of T correspond exactly to the vertical line segments in the symbol square which lie above them, but after the first amalgamation has taken place this is no longer the case. The dotted lines shown in the figures, which join points in the symbol square to points of the tree, are therefore only suggestive.

Example 13 Consider the pruning disk $D = D(\bar{0}101 * \bar{1}\bar{0})$ of Example 3. In this case y contains two symbols after $\bar{0}1$, so $n = 3$ and it is necessary to amalgamate $F^{-1}(D)$ and $F^{-2}(D)$ with \mathbb{T} . The amalgamation of $F^{-1}(D)$ is exactly as shown in Figure 47, with $y = \bar{0}101$ and $x = \bar{1}\bar{0}$. This




Figure 47: Amalgamating $F^{-1}(D)$ with \mathbb{T}

yields a tree map $f_1: T_1 \rightarrow T_1$, where T_1 is a triod. The effect on f_1 of amalgamating $F^{-2}(D)$ with $\mathbb{T} \cup F^{-1}(D)$ is shown in Figure 48: pairs of points with itineraries $\cdot 11z$ and $\cdot 10z$, where $z \succeq \overline{10}$, are identified (the two points in such a pair have the same image under f_1 , since pairs of the form $\cdot 0z$ and $\cdot 1z$ have already been identified). This yields a new tree map $f_2: T_2 \rightarrow T_2$, where T_2 has two vertices of valence three. Figure 49 shows that $F^{-3}(D)$ corresponds to a segment K of T_2 with the property that $f_2(K)$ backtracks: pruning away $F^{-3}(D)$ corresponds to pulling K tight, yielding the graph map $f_D: T_2 \rightarrow T_2$ corresponding to the pruned horseshoe F_D .




Figure 48: Amalgamating $F^{-2}(D)$ with $\mathbb{T} \cup F^{-1}(D)$

Example 14 Consider the pruning front consisting of the two pruning disks $D_1 = D(\overline{010010} * \overline{100})$ and $D_2 = D(\overline{011} * \overline{1})$ (see Examples 4 and 7). We start with D_1 for which $n = 5$, so it is necessary to amalgamate $F^{-1}(D_1)$ through $F^{-4}(D_1)$ with \mathbb{T} . Amalgamating $F^{-1}(D_1)$ proceeds exactly as in Example 13 (i.e. as shown in Figure 47): the resulting tree map $f_1: T_1 \rightarrow T_1$ is shown in Figure 50.

The second preimage $F^{-2}(D_1)$ contains a complete vertical strip in the right half of the symbol square in addition to the smaller rectangle in the left half (Figure 51): this is reflected by the fact that there are two segments K_0 and K_1 of T_1 whose images under f_1 backtrack. In such




Figure 49: Pruning away $F^{-3}(D)$




Figure 50: The quotient tree map after amalgamating $F^{-1}(D_1)$ with \mathbb{T}

a situation, it is necessary to begin by pruning away the part of $F^{-2}(D_1)$ corresponding to the complete vertical strip, or, equivalently, by pulling tight K_1 as shown in Figure 51 to give the tree map $f'_1: T_1 \rightarrow T_1$. The part of D_1 corresponding to the complete vertical strip in $F^{-2}(D_1)$ has now been pruned away, and we need only consider the remainder of the disk.




Figure 51: Pruning away a complete vertical strip

Amalgamating the remainder of $F^{-2}(D_1)$ corresponds to identifying pairs of points with itineraries $\cdot 00z$ and $\cdot 01z$, where $z \succeq \overline{100}$, to give the tree map $f_2: T_2 \rightarrow T_2$ shown in Figure 52.

The amalgamation of the remainder of $F^{-3}(D_1)$ and $F^{-4}(D_1)$ is shown in Figure 53: on the level of the tree maps, these correspond to identifying pairs of points with itineraries $(\cdot 101w, \cdot 100w)$ and $(\cdot 0101w, \cdot 0100w)$ respectively, where $w \succeq \overline{100}$. We are left with a tree map $f_4: T_4 \rightarrow T_4$, where T_4 has a segment K corresponding to (the remains of) $F^{-5}(D_1)$ whose image backtracks: pulling K tight yields the tree map $f_{D_1}: T_4 \rightarrow T_4$ corresponding to the pruned horseshoe F_{D_1} (Figure 54).

Finally, we deal with D_2 . Since $n = 2$ in this case, only a single amalgamation is necessary. This amalgamation identifies pairs of points $(\cdot 0w, \cdot 1w)$, where $w \succeq \overline{1}$ (D_2 is the same disk considered in Example 2). Pulling tight a backtracking segment corresponding to $F^{-2}(D_2)$ gives the tree map $f_{D_1 D_2}$ corresponding to the pruned horseshoe $F_{D_1 D_2}$ (Figure 55).

B Appendix: From tree maps to horseshoe symbolics

When the tree map corresponding to a pruned horseshoe is constructed as in Appendix A, it is straightforward to find a Markov graph describing the orbits which survive the pruning in terms of the original symbolics of the horseshoe. All that is necessary is to keep track of which edges of the quotient tree come from each of the two halves of the initial interval I . Because the first




Figure 52: Amalgamating (the remains of) $F^{-2}(D_1)$

amalgamation (shown in Figure 47) always identifies pairs of points symmetrically placed about the centre of I , it is also necessary to keep track of both sides of this first new edge: one side corresponds to the symbol 0, and the other to the symbol 1.

Figure 56 shows the tree map corresponding to the pruned horseshoe F_D , where $D = D(\overline{0101} * \overline{10})$ (see Example 13). The two sides of the first new edge are labelled b and \hat{b} : since the edge d has image b , it is also necessary to distinguish between the two sides d and \hat{d} of this edge. Following through the construction of Figures 47–49, it can easily be seen that the edges a and b come from the left hand side of I , and hence correspond to the symbol 0, while the other edges correspond to the symbol 1. We can therefore construct a Markov graph (also shown in Figure 56) which describes, in terms of the original symbolics, the horseshoe orbits which survive pruning (up to the usual ambiguity arising from intersections of edges at their endpoints).

References

- [1] J. Alves and J. Sousa Ramos. Kneading theory for tree maps. To appear in *Ergodic Theory Dynam. Systems*, 2001.
- [2] R. Artuso, E. Aurell, and P. Cvitanović. Recycling of strange sets. I. Cycle expansions. *Nonlinearity*, 3(2):325–359, 1990.
- [3] R. Artuso, E. Aurell, and P. Cvitanović. Recycling of strange sets. II. Applications. *Nonlinearity*, 3(2):361–386, 1990.
- [4] M. Baillif. Dynamical zeta functions for tree maps. *Nonlinearity*, 12(6):1511–1529, 1999.




Figure 53: Amalgamating (the remains of) $F^{-3}(D_1)$ and $F^{-4}(D_1)$




Figure 54: The tree map corresponding to F_{D_1}

- [5] M. Baillif and A. de Carvalho. Piecewise linear model for tree maps. *Internat. J. Bifur. Chaos Appl. Sci. Engrg.*, 11(12):3163–3169, 2001.
- [6] M. Barge and J. Martin. The construction of global attractors. *Proc. Amer. Math. Soc.*, 110(2):523–525, 1990.
- [7] E. Bedford, M. Lyubich, and J. Smillie. Distribution of periodic points of polynomial diffeomorphisms of \mathbf{C}^2 . *Invent. Math.*, 114(2):277–288, 1993.
- [8] E. Bedford, M. Lyubich, and J. Smillie. Polynomial diffeomorphisms of \mathbf{C}^2 . IV. The measure of maximal entropy and laminar currents. *Invent. Math.*, 112(1):77–125, 1993.
- [9] E. Bedford and J. Smillie. Fatou-Bieberbach domains arising from polynomial automorphisms. *Indiana Univ. Math. J.*, 40(2):789–792, 1991.
- [10] E. Bedford and J. Smillie. Polynomial diffeomorphisms of \mathbf{C}^2 : currents, equilibrium measure and hyperbolicity. *Invent. Math.*, 103(1):69–99, 1991.
- [11] E. Bedford and J. Smillie. Polynomial diffeomorphisms of \mathbf{C}^2 . II. Stable manifolds and recurrence. *J. Amer. Math. Soc.*, 4(4):657–679, 1991.
- [12] E. Bedford and J. Smillie. Polynomial diffeomorphisms of \mathbf{C}^2 . III. Ergodicity, exponents and entropy of the equilibrium measure. *Math. Ann.*, 294(3):395–420, 1992.




Figure 55: Amalgamating $F^{-1}(D_2)$ and pruning it away




Figure 56: Horseshoe symbolics for the orbits which survive pruning

- [13] E. Bedford and J. Smillie. Polynomial diffeomorphisms of \mathbf{C}^2 . V. Critical points and Lyapunov exponents. *J. Geom. Anal.*, 8(3):349–383, 1998.
- [14] E. Bedford and J. Smillie. Polynomial diffeomorphisms of \mathbf{C}^2 . VI. Connectivity of J . *Ann. of Math. (2)*, 148(2):695–735, 1998.
- [15] E. Bedford and J. Smillie. External rays in the dynamics of polynomial automorphisms of \mathbf{C}^2 . In *Complex geometric analysis in Pohang (1997)*, pages 41–79. Amer. Math. Soc., Providence, RI, 1999.
- [16] E. Bedford and J. Smillie. Polynomial diffeomorphisms of \mathbf{C}^2 . VII. Hyperbolicity and external rays. *Ann. Sci. École Norm. Sup. (4)*, 32(4):455–497, 1999.
- [17] M. Benedicks and L. Carleson. On the Hénon attractor. In *IXth International Congress on Mathematical Physics (Swansea, 1988)*, pages 498–500. Hilger, Bristol, 1989.
- [18] M. Benedicks and L. Carleson. The dynamics of the Hénon map. *Ann. of Math. (2)*, 133(1):73–169, 1991.
- [19] M. Bestvina and M. Handel. Train-tracks for surface homeomorphisms. *Topology*, 34(1):109–140, 1995.
- [20] O. Biham and W. Wenzel. Characterization of unstable periodic orbits in chaotic attractors and repellers. *Phys. Rev. Lett.*, 63(8):819–822, 1989.
- [21] R. Bowen. Entropy for group endomorphisms and homogeneous spaces. *Trans. Amer. Math. Soc.*, 153:401–414, 1971.
- [22] P. Boyland. Braid types and a topological method of proving positive entropy. Preprint, Boston University, 1984.
- [23] P. Boyland. Topological methods in surface dynamics. *Topology Appl.*, 58(3):223–298, 1994.

- [24] P. Collins. Dynamics forced by surface trellises. In *Geometry and topology in dynamics (Winston-Salem, NC, 1998/San Antonio, TX, 1999)*, pages 65–86. Amer. Math. Soc., Providence, RI, 1999.
- [25] P. Cvitanović. Invariant measurement of strange sets in terms of cycles. *Phys. Rev. Lett.*, 61(24):2729–2732, 1988.
- [26] P. Cvitanović. Periodic orbits as the skeleton of classical and quantum chaos. *Phys. D*, 51(1-3):138–151, 1991. Nonlinear science: the next decade (Los Alamos, NM, 1990).
- [27] P. Cvitanović. Periodic orbit theory in classical and quantum mechanics. *Chaos*, 2(1):1–4, 1992.
- [28] P. Cvitanović, G. Gunaratne, and I. Procaccia. Topological and metric properties of Hénon-type strange attractors. *Phys. Rev. A (3)*, 38(3):1503–1520, 1988.
- [29] P. Dahlqvist. On the effect of pruning on the singularity structure of zeta functions. *J. Math. Phys.*, 38(8):4273–4282, 1997.
- [30] M. J. Davis, R. S. MacKay, and A. Sannami. Markov shifts in the Hénon family. *Phys. D*, 52(2-3):171–178, 1991.
- [31] A. de Carvalho. Extensions, quotients and generalized pseudo-anosov maps. To appear in *Graphs and Patterns in Mathematics and Theoretical Physics*.
- [32] A. de Carvalho. Pruning fronts and the formation of horseshoes. *Ergodic Theory Dynam. Systems*, 19(4):851–894, 1999.
- [33] A. de Carvalho and T. Hall. Braid forcing and star-shaped train tracks. Submitted.
- [34] A. de Carvalho and T. Hall. Decoration invariants for horseshoe braids. In preparation.
- [35] A. de Carvalho and T. Hall. The forcing relation for horseshoe braid types. To appear in *Experiment. Math.*, 2001.
- [36] A. de Carvalho and T. Hall. Pruning theory and Thurston’s classification of surface homeomorphisms. *J. Eur. Math. Soc. (JEMS)*, 3(4):287–333, 2001.
- [37] A. de Carvalho and M. Paternain. Monotone quotients of surface diffeomorphisms. In preparation.
- [38] R. Devaney. *An introduction to chaotic dynamical systems*. Addison-Wesley Publishing Company Advanced Book Program, Redwood City, CA, second edition, 1989.
- [39] R. Devaney and Z. Nitecki. Shift automorphisms in the Hénon mapping. *Comm. Math. Phys.*, 67(2):137–146, 1979.
- [40] H. El Hamouly and C. Mira. Lien entre les propriétés d’un endomorphisme de dimension un et celles d’un difféomorphisme de dimension deux. *C. R. Acad. Sci. Paris Sér. I Math.*, 293(10):525–528, 1981.

- [41] A. Fathi, F. Laudenbach, and V. Poénaru. *Travaux de Thurston sur les surfaces*. Société Mathématique de France, Paris, 1979. Séminaire Orsay.
- [42] S. Fenley. End periodic surface homeomorphisms and 3-manifolds. *Math. Z.*, 224(1):1–24, 1997.
- [43] J. Franks and M. Misiurewicz. Cycles for disk homeomorphisms and thick trees. In *Nielsen theory and dynamical systems (South Hadley, MA, 1992)*, pages 69–139. Amer. Math. Soc., Providence, RI, 1993.
- [44] J.-M. Gambaudo and C. Tresser. Simple models for bifurcations creating horseshoes. *J. Statist. Phys.*, 32(3):455–476, 1983.
- [45] P. Grassberger and H. Kantz. Generating partitions for the dissipative Hénon map. *Phys. Lett. A*, 113(5):235–238, 1985.
- [46] P. Grassberger, H. Kantz, and U. Moenig. On the symbolic dynamics of the Hénon map. *J. Phys. A*, 22(24):5217–5230, 1989.
- [47] T. Hall. Prune: Software available from <http://www.liv.ac.uk/~tobyhall/prune/>.
- [48] T. Hall. The creation of horseshoes. *Nonlinearity*, 7(3):861–924, 1994.
- [49] T. Hall. Fat one-dimensional representatives of pseudo-Anosov isotopy classes with minimal periodic orbit structure. *Nonlinearity*, 7(2):367–384, 1994.
- [50] M. Handel. Private communication.
- [51] M. Handel. A fixed-point theorem for planar homeomorphisms. *Topology*, 38(2):235–264, 1999.
- [52] K. Hansen. Pruning of orbits in four-disk and hyperbola billiards. *Chaos*, 2(1):71–75, 1992.
- [53] K. Hansen. Remarks on the symbolic dynamics for the Hénon map. *Phys. Lett. A*, 165(2):100–104, 1992.
- [54] K. Hansen. Symbolic dynamics. I. Finite dispersive billiards. *Nonlinearity*, 6(5):753–769, 1993.
- [55] K. Hansen. *Symbolic dynamics in chaotic systems*. PhD thesis, University of Oslo, 1993.
- [56] K. Hansen and P. Cvitanović. Bifurcation structures in maps of Hénon type. *Nonlinearity*, 11(5):1233–1261, 1998.
- [57] M. Hénon. A two-dimensional mapping with a strange attractor. *Comm. Math. Phys.*, 50(1):69–77, 1976.
- [58] P. Holmes and D. Whitley. Bifurcations of one- and two-dimensional maps. *Philos. Trans. Roy. Soc. London Ser. A*, 311(1515):43–102, 1984.
- [59] P. Holmes and R. Williams. Knotted periodic orbits in suspensions of Smale’s horseshoe: torus knots and bifurcation sequences. *Arch. Rational Mech. Anal.*, 90(2):115–194, 1985.

- [60] J. Hubbard and R. Oberste-Vorth. Hénon mappings in the complex domain. I. The global topology of dynamical space. *Inst. Hautes Études Sci. Publ. Math.*, (79):5–46, 1994.
- [61] J. Hubbard and R. Oberste-Vorth. Hénon mappings in the complex domain. II. Projective and inductive limits of polynomials. In *Real and complex dynamical systems (Hillerød, 1993)*, pages 89–132. Kluwer Acad. Publ., Dordrecht, 1995.
- [62] H. Hulme. *Finite and infinite braids: a dynamical systems approach*. PhD thesis, University of Liverpool, 2000.
- [63] Y. Ishii. Towards a kneading theory for Lozi mappings. I. A solution of the pruning front conjecture and the first tangency problem. *Nonlinearity*, 10(3):731–747, 1997.
- [64] Y. Ishii. Towards a kneading theory for Lozi mappings. II. Monotonicity of the topological entropy and Hausdorff dimension of attractors. *Comm. Math. Phys.*, 190(2):375–394, 1997.
- [65] Y. Ishii and D. Sands. Monotonicity of the Lozi family near the tent-maps. *Comm. Math. Phys.*, 198(2):397–406, 1998.
- [66] I. Kan, H. Koçak, and J. Yorke. Antimonotonicity: concurrent creation and annihilation of periodic orbits. *Ann. of Math. (2)*, 136(2):219–252, 1992.
- [67] A. Katok. Lyapunov exponents, entropy and periodic orbits for diffeomorphisms. *Inst. Hautes Études Sci. Publ. Math.*, (51):137–173, 1980.
- [68] A. Katok and B. Hasselblatt. *Introduction to the modern theory of dynamical systems*. Cambridge University Press, Cambridge, 1995. With a supplementary chapter by Katok and Leonardo Mendoza.
- [69] J. Los. Pseudo-Anosov maps and invariant train tracks in the disc: a finite algorithm. *Proc. London Math. Soc. (3)*, 66(2):400–430, 1993.
- [70] R. Lozi. Strange attractors: a class of mappings of \mathbb{R}^2 which leaves some Cantor sets invariant. In *Intrinsic stochasticity in plasmas (Internat. Workshop, Inst. Études Sci. Cargèse, Cargèse, 1979)*, pages 373–381. École Polytech., Palaiseau, 1979.
- [71] J. Milnor and W. Thurston. On iterated maps of the interval. In *Dynamical systems (College Park, MD, 1986–87)*, pages 465–563. Springer, Berlin, 1988.
- [72] R. Moore. *Foundations of point set theory*. American Mathematical Society, Providence, R.I., 1962.
- [73] J. Roberts and N. Steenrod. Monotone transformations of two-dimensional manifolds. *Ann. of Math.*, 39:851–862, 1938.
- [74] S. Smale. Differentiable dynamical systems. *Bull. Amer. Math. Soc.*, 73:747–817, 1967.
- [75] D. Sterling, H. R. Dullin, and J. D. Meiss. Homoclinic bifurcations for the Hénon map. *Phys. D*, 134(2):153–184, 1999.

- [76] W. Thurston. On the geometry and dynamics of diffeomorphisms of surfaces. *Bull. Amer. Math. Soc. (N.S.)*, 19(2):417–431, 1988.
- [77] H. C. Tseng, H. J. Chen, P. C. Li, W. Y. Lai, C. H. Chou, and H. W. Chen. Some exact results for the diffusion coefficients of maps with pruning cycles. *Phys. Lett. A*, 195(1):74–80, 1994.
- [78] Q. Wang and L-S. Young. Strange attractors with one direction of instability. *Comm. Math. Phys.*, 218(1):1–97, 2001.