

Topological Entropy Bounds for Hyperbolic Plateaus of the Hénon Map

Rafael M. Frongillo*

Advisor: John Smillie

May 15, 2014

Abstract

Combining two existing rigorous computational methods, for verifying hyperbolicity (due to Arai) and for computing topological entropy bounds (due to Day et al.), we prove lower bounds on topological entropy for 43 hyperbolic plateaus of the Hénon map. We also examine the 16 area-preserving plateaus studied by Arai and compare our results with related work. Along the way, we augment the algorithms of Day et al. with routines to optimize the algorithmic parameters and simplify the resulting semi-conjugate subshift.

1 Introduction

Dynamical systems theory has seen the emergence of many rigorous computational methods in recent years. Such tools often extend the realm of provable theorems well beyond what is possible with chalk and blackboard. This is particularly true of the recent automated tools to compute topological entropy bounds [18, 6] and to prove hyperbolicity [1, 13, 16].

Let us briefly recall the relevant characteristics of these methods. Both techniques for proving entropy bounds first construct a subshift of finite type (SFT), whose topological entropy is easily computed and is a lower bound of that of the original system. Newhouse et al. compute rigorous approximations of stable and unstable manifolds of periodic orbits, and then construct a SFT using pieces of these manifolds; in this regard, their technique could be considered a rigorous version of the trellis method. Day, Frongillo, and Treviño (DFT) [6] construct a discrete multivalued map from a discretization of the phase space, and then apply discrete Conley index theory to prove a semi-conjugacy to a particular SFT. In certain settings, such as the one we study here, the DFT method is *completely* automated, meaning that after a simple initialization, no further manual input is required. It is unclear whether the method of Newhouse et al. shares this automation property.

Hruska [13] developed one of the first automated methods for rigorously verifying hyperbolicity, based on the computation of cone fields. In contrast, Arai [1] employs a

*raf@cs.berkeley.edu. Work done at Cornell University and the University of California at Berkeley.

more indirect technique using the notion of quasi-hyperbolicity. His approach allows for more efficient computations than Hruska, but does not guarantee that the nonwandering set is not just a finite collection of periodic orbits, or even that it is nonempty.

In [1], Arai identifies several hyperbolic regions of the Hénon map:

$$f_{a,b}(x, y) = (a - x^2 + by, x) \tag{1}$$

These regions are dubbed *hyperbolic plateaus* because topological entropy is constant across any such region. Arai’s technique does not reveal anything about these topological entropy values, however, and it is therefore natural to combine his computations with an automated method for proving topological entropy bounds.

In this paper we use the DFT method [6] to compute lower bounds for topological entropy of the hyperbolic plateaus of Hénon computed by Arai in [1]. The constant entropy on each plateau enables us to extend a lower bound computation from a single setting of parameter values to an entire region of the parameter space. Additionally, the full automation of the DFT method enables us to study a total of 58 parameter values with essentially the same manual effort as studying one. We use the DFT method as a black box, but we give a high-level overview of the approach in Section 4, including new techniques for improving robustness and simplifying the resulting SFT.

Theorems 5.1 and 5.3 summarize our results. To the author’s knowledge, all of these rigorous lower bounds are the largest known for their corresponding parameters. We selected the parameter regions so that the bounds obtained might give a global picture of the entropy of the Hénon map as a function of the parameters; see Figure 1 for such a picture.

We also study several of the area-preserving Hénon maps in Section 5.1, which have been well-studied in the Physics community [12, 23, 7]. Recently, some precise rigorous results emerged as well [2]. We find that our rigorous lower bounds match or are very close to estimates given in previous work, and match the rigorous results exactly when applicable.

2 Background

We first review basic definitions related to symbolic dynamics, topological entropy, and hyperbolicity, in Sections 2.1 and 2.2. The DFT approach relies heavily on a combinatorial version of the discrete Conley index; in Section 2.3 we discuss the index at a high level, and introduce the combinatorial structures that relate it to our computational setting.

2.1 Symbolic dynamics and topological entropy

Define a *symbol space* $X_n = \{0, \dots, n - 1\}^{\mathbb{Z}}$ to be the set of all bi-infinite sequences on n symbols. It is well-known that X_n is a complete metric space. Let the *full n -shift* $\sigma : X_n \rightarrow X_n$ be the map acting on X_n by $(\sigma(x))_i = x_{i+1}$. Given a directed graph G on n nodes with $n \times n$ transition matrix A with $A_{i,j} \in \{0, 1\}$, we can define an induced symbol space $X_G \subseteq X_n$, where $x \in X_G$ if and only if for $x = (\dots, x_i, x_{i+1}, \dots)$, $A_{x_i, x_{i+1}} = 1$ for all i . That is, X_G consists of all sequences in X_n with transitions of σ allowed by the




Figure 1: Rigorous lower bounds for topological entropy for the hyperbolic plateaus in Figure 4. The height of each plateau in the visualization is proportional to the entropy bound computed. See Theorem 5.1 or Figure 5 for the actual bounds.

edges of G . Equipped with the corresponding shift map $\sigma_G : X_G \rightarrow X_G$, we call (X_G, σ_G) a *subshift of finite type*.

We use topological entropy to measure the relative complexity of different dynamical systems. If the topological entropy of a dynamical system f , denoted $h(f)$, is positive, we say that f is *chaotic*.

Definition 2.1 ([15], adapted from [3]). *Let $f : X \rightarrow X$ be a continuous map with respect to a metric d . We say that a set $W \subseteq X$ is (n, ε) -separated under f if for any distinct $x, y \in W$ we have $d(f^j(x), f^j(y)) > \varepsilon$ for some $0 \leq j < n$. The topological entropy of f is*

$$h(f) := \lim_{\varepsilon \rightarrow \infty} \limsup_{n \rightarrow \infty} \frac{\log(s_f(n, \varepsilon))}{n}, \quad (2)$$

where $s_f(n, \varepsilon)$ denotes the maximum cardinality of an (n, ε) -separated set under f .

While topological entropy can be difficult to calculate in general, there is a simple formula for subshifts of finite type which is given in the following theorem. For a proof,

see [15] or [21].

Theorem 2.2. *Let G be a directed graph with transition matrix A , and let (X_G, σ_G) be the corresponding subshift of finite type. Then the topological entropy of σ_G is $h(\sigma_G) = \log(\text{sp}(A))$, where $\text{sp}(A)$ denotes the spectral radius (maximum magnitude of an eigenvalue) of A .*

When studying a complex map f , it is sometimes useful to study a subsystem of f which can be precisely related to f via a semi-conjugacy.

Definition 2.3. *Let $f : X \rightarrow X$ and $g : Y \rightarrow Y$ be continuous maps for topological spaces X, Y . A semi-conjugacy from f to g is a continuous surjection $\phi : X \rightarrow Y$ with $\phi \circ f = g \circ \phi$. We say that f is semi-conjugate to g if there exists a semi-conjugacy from f to g . If additionally ϕ is a homeomorphism, then f and g are conjugate.*

Particularly relevant to our setting is the following result.

Theorem 2.4 ([21]). *Let f and g be continuous maps, and let ϕ be a semi-conjugacy from f to g . Then $h(f) \geq h(g)$.*

Note that if f and g are conjugate, Theorem 2.4 gives us $h(f) = h(g)$. In other words, topological entropy is invariant under conjugacy.

2.2 Hyperbolicity

To begin we define uniform hyperbolicity. Throughout the paper, we will refer to this property simply as hyperbolicity. Let X be a manifold.

Definition 2.5 ([11]). *A map $f : X \rightarrow X$ is said to be (uniformly) hyperbolic if for every $x \in X$ the tangent space $T_x X$ for f is a direct sum of stable and unstable subspaces; more precisely, if we have $T_x X = E^s(x) \oplus E^u(x)$, where $E^s(x)$ and $E^u(x)$ satisfy the following inequalities for some $C > 0$ and $0 < \lambda < 1$, and for all $n \in \mathbb{N}$:*

1. $\|Df^n(v)\| \leq C\lambda^n \|v\|$ for all $v \in E^s(x)$.
2. $\|Df^{-n}(v)\| \leq C\lambda^n \|v\|$ for all $v \in E^u(x)$.

This structure can be thought of as a generalization of the structure of the Smale horseshoe, namely that there are invariant directions, and there is uniform contraction and expansion in the stable and unstable directions, respectively. Some useful properties of hyperbolic systems are discussed below, but for more details see [21] and [11].

An important property of hyperbolic maps is that they are structurally stable [21], which implies that all maps in the same hyperbolic region are conjugate. Thus, by Theorem 2.4, the topological entropy is constant within such a region. For this reason, we will henceforth call these regions *hyperbolic plateaus*.

Hyperbolicity often makes it easier to identify interesting dynamics, but it is important to note that sometimes a system can be “vacuously” hyperbolic, in the sense that it is hyperbolic but there is no recurrent behavior. A helpful concept in this context is the *nonwandering set*.

Definition 2.6 ([21]). *The nonwandering set of a map f is the set of points x for which every neighborhood U of x has $f^n(U) \cap U \neq \emptyset$ for some $n \geq 1$.*

2.3 Combinatorial structures and the discrete Conley index

As we will be performing rigorous computations, we will naturally be interested in finite representations of the continuous map $f : X \rightarrow X$ in question. Our main tool in this regard is a *combinatorial enclosure*, a discrete “outer-approximation” of f . We first discuss how we discretize the phase space X , which we assume here to be a subset of \mathbb{R}^n , though the constructions that follow generalize to more complicated spaces as well.

We begin by setting up a grid \mathcal{G} on X , composed of finitely many cubical complices \mathcal{B}_i . In practice, all elements of the grid are rectangles, which we call *boxes*, represented as products of intervals (viewed in some nice coordinate chart); that is, for each $\mathcal{B}_i \in \mathcal{G}$, we can write $\mathcal{B}_i = \prod_{k=1}^n [x_k^i, y_k^i]$ for some $x^i, y^i \in \mathbb{R}^n$. In the present paper, for each k the interval widths $y_k^i - x_k^i$ are constant for all i , meaning the shape of all boxes is the same. For a collection of boxes $\mathcal{K} \subseteq \mathcal{G}$, we denote by $|\mathcal{K}|$ its *topological realization*, that is, its corresponding subset of X .

In our setting, we create \mathcal{G} by selecting one box \mathcal{B} , called the *initial box*, which encloses the entire area we wish to study. We then subdivide \mathcal{B} evenly d times in each coordinate direction, meaning we replace $|\mathcal{B}|$ with 2^n boxes by splitting each coordinate in half, and then we recurse on each of those boxes, with total recursion depth d . This gives a grid \mathcal{G} with 2^{dn} boxes, each of size 2^{-d} relative to the initial box \mathcal{B} . The integer d will be referred to as the *resolution* or *depth*.

After discretizing the phase space, we must discretize the map f itself. To do this we introduce the combinatorial enclosure.

Definition 2.7. A multivalued map $F : X \rightrightarrows X$ is a map from X to its power set, so that $F(x) \subseteq X$. If F is acyclic¹ and we have $f(x) \in F(x) \forall x \in X$ for some continuous single-valued map f , then F is an enclosure of f .

Definition 2.8. Given an enclosure F of f , a combinatorial enclosure of f is a multivalued map $\mathcal{F} : \mathcal{G} \rightrightarrows \mathcal{G}$ defined by $\mathcal{F}(\mathcal{B}) = \{\mathcal{B}' \in \mathcal{G} : |\mathcal{B}'| \cap F(|\mathcal{B}|) \neq \emptyset\}$.

The reason multivalued maps and enclosures are used in our computations is that, if constructed properly, they enable rigorous results. In particular, for each $\mathcal{B} \in \mathcal{G}$ we use the interval arithmetic library Intlab [22] to compute rigorous outer approximations $Y_{\mathcal{B}}$ of $f(|\mathcal{B}|)$, and let $F(x) = Y_{\mathcal{B}}$ for all $x \in |\mathcal{B}|$. Defining \mathcal{F} as in Definition 2.8, this technique keeps track of the error terms in the computations of the image of a box, ensuring that the true image $f(|\mathcal{B}|)$ is contained in $|\mathcal{F}(\mathcal{B})|$. Note that $|\mathcal{F}|$ is also an enclosure of f .

In this combinatorial setting, we can apply a version of the discrete Conley index to gather rigorous information about the dynamics of the underlying map. We will not define the Conley index formally here, giving instead a high-level overview of the concepts immediately relevant to our study. For a full treatment of Conley index theory, see [20, 17]; for the definitions and algorithms needed in our discrete setting, see [6].

Given some compact set $N \subseteq X$, its *maximal invariant set* is the set of points $x \in N$ such that there exists a trajectory of f through x which stays entirely within N for all forward and backward time.² A central concept of Conley index theory is that of *isolation*:

¹ F is acyclic if for all x , $F(x)$ has trivial homology on all levels (no “holes”).

²That is, there exists some $\{x_n\}_{n \in \mathbb{Z}} \subseteq N$ such that $f(x_n) = x_{n+1}$ for all $n \in \mathbb{Z}$ and $x_0 = x$.

we say that an invariant set S (that is, a set with $f(S) = S$) is *isolated* if there is a compact set N , called the *isolating neighborhood* of S , such that S is contained in the interior of N and S is the maximal invariant set of N . From this notion, a (*combinatorial*) *index pair* for the enclosure \mathcal{F} is a pair $(\mathcal{P}_1, \mathcal{P}_0)$ of sets of boxes satisfying a two properties: (a) letting $P_i = |\mathcal{P}_i|$, the set $\overline{P_1} \setminus P_0$ must be an isolating neighborhood, and (b) if we contract the set P_0 to a point denoted $[P_0]$, the map f induces on the pointed quotient space $(P_1/P_0, [P_0])$ must be continuous.³ Finally, the *Conley index* of the index pair $(\mathcal{P}_1, \mathcal{P}_0)$ is a particular topological invariant of its maximal invariant set based on relative homology; in our setting, the index can be represented as a finite sequence of matrices describing the map that f induces on each level of homology.⁴

With the Conley index, we can make strong statements about the behavior of f . For example, if the Conley index corresponding to index pair $(\mathcal{P}_1, \mathcal{P}_0)$ is nontrivial,⁵ the maximal invariant set of $\overline{P_1} \setminus P_0$ is nonempty. Along the same lines, but using a more nuanced analysis, the DFT method is able to take the Conley index and compute a SFT which is semi-conjugate to f . We describe the DFT method at a high level in Section 4.1. Note however that computing the Conley index in the form of an induced map on homology is a difficult task in and of itself. For a thorough treatment of computational homology, dealing with this and other applications, see [14]. For our computations, we use the *homcubes* package [19], part of the software package *CHomP* [4].

3 Simplifying Subshifts

Given a subshift of finite type (X_G, σ_G) for a graph G , it is often of interest to know whether there is a graph H on fewer vertices such that (X_G, σ_G) and (X_H, σ_H) are conjugate. To this end, we recall the notion of strong shift equivalence.

Definition 3.1 (Strong shift equivalence). *Let A and B be matrices. An elementary shift equivalence between A and B is a pair (R, S) such that*

$$A = RS \text{ and } B = SR. \tag{3}$$

In this case, we write $(R, S) : A \rightarrow B$. If there is a sequence of such elementary shift equivalences $(R_i, S_i) : A_{i-1} \rightarrow A_i$, $1 \leq i \leq k$, we say that A_0 and A_k are strongly shift equivalent.

This notion is useful because of the following result due to R. F. Williams.

Theorem 3.2 ([24]). *For directed graphs G and H , the corresponding subshifts (X_G, σ_G) and (X_H, σ_H) are conjugate if and only if the transition matrices of G and H are strongly shift equivalent.*

³We define the induced map g on this space as follows: define $g([P_0]) = [P_0]$, and for $x \in P_1 \setminus P_0$, define $g(x) = [P_0]$ if $f(x) \in P_0$ and $g(x) = f(x)$ otherwise.

⁴Specifically, the index is the shift equivalence class of the induced map on relative homology groups of the quotient space. Again, for the full details, see [20, 17, 6].

⁵The index is trivial if it is the shift equivalence class of the zero map.

Theorem 3.2 allows us to prove that two subshifts are conjugate by a series of simple matrix computations. Finding matrices that give a strong shift equivalence, however, can be a very difficult problem. Two methods of finding such equivalences are given in [15]: state splitting, where a single vertex is split into two, or state amalgamation, where two vertices are combined into one. In graph-theoretic terms, amalgamating two vertices is equivalent to contracting them together. In general obtaining the smallest element of a strong shift equivalence class may involve both splittings and amalgamations. We instead focus on the simpler problem of obtaining H only by amalgamating vertices in G . This also has the advantage of producing a matrix that is more useful for our needs in this paper (see Section 4).

Let A be the binary $n \times n$ transition matrix for G . The following two conditions, adapted from [15], will allow us to amalgamate vertices i and j .

$$\textbf{Forward Condition: } A\vec{e}_i = A\vec{e}_j \text{ and } (\vec{e}_i^\top A) \cdot (\vec{e}_j^\top A) = 0 \quad (4)$$

$$\textbf{Backward Condition: } \vec{e}_i^\top A = \vec{e}_j^\top A \text{ and } (A\vec{e}_i) \cdot (A\vec{e}_j) = 0 \quad (5)$$

Here \vec{e}_i denotes the column vector with a 1 in position i and zeros elsewhere. From a graph-theoretic or dynamical systems point of view, the forward condition says that i and j have the same image but disjoint preimages, and the backward condition says they have the same preimage but disjoint images. See Figure 2 for an example.




Figure 2: A forward amalgamation

Note that the backward condition for A is the same as the forward condition for A^\top . The following result allows us to reduce A to a smaller $n - 1$ by $n - 1$ matrix B if either of these conditions are satisfied for some pair of vertices. See e.g. [15, §2] for a proof.

Theorem 3.3. *If i and j satisfy the forward condition (4) or backward condition (5) for a transition matrix A , then there is an elementary shift equivalence from A to the matrix obtained by amalgamating i and j .*

By applying Theorem 3.3 repeatedly, as long as there exist i, j satisfying either the forward or backward condition, one can reduce A to a much smaller representative of its strong shift equivalence class. The resulting matrix B at the end of this process corresponds to a subshift (X_H, σ_H) which is therefore conjugate to (X_G, σ_G) .

For small enough matrices it is feasible to perform a simple brute-force search to find the smallest B which can be obtained from A via amalgamations, but we would like a more efficient algorithm for larger matrices. Unfortunately, it is shown in [10] that it is NP-hard (computationally intractable) to find an ordering of amalgamations which yields the smallest representative. In light of this hardness, we use the procedure outlined in

Algorithm 1 `simplify_subshift`: Amalgamating a subshift of finite type

```
Input: subshift  $T \in \{0,1\}^{n \times n}$ , number of trials  $K$ 
 $T_{\min} \leftarrow T$ 
for  $k$  from 1 to  $K$  do
   $\pi \leftarrow \text{random\_permutation}(n)$ 
   $T_\pi \leftarrow T(\pi, \pi)$  {relabel vertices}
  repeat
     $\text{amalgamated} \leftarrow \text{false}$ 
    for  $(i, j) \in E(T_\pi)$  ordered lexicographically do
      if conditions (4) or (5) hold for  $A = T_\pi$  then
         $T_\pi \leftarrow \text{amalgamate}(T_\pi, i, j)$ 
         $\text{amalgamated} \leftarrow \text{true}$ 
      break for
    end if
  end for
  until not  $\text{amalgamated}$  {no further amalgamations}
  if  $\text{size}(T_\pi) < \text{size}(T_{\min})$  then
     $T_{\min} \leftarrow T_\pi$ 
  end if
end for
Output:  $T_{\min}$ 
```

Algorithm 1, which is essentially a randomized greedy algorithm performed many times. Taking the number of trials k to be about n^2 will typically give a reasonable approximation factor, meaning that if m amalgamations are possible, the algorithm will perform roughly $m/3$ or $m/2$ amalgamations. It remains an open question whether an algorithm exists which has a provable approximation guarantee.

4 Techniques

Given a continuous map $f : X \rightarrow X$, we will apply discrete Conley index theory to compute subshifts of finite type to which f is semi-conjugate. This is the approach of the DFT method [6], which we describe first in Section 4.1. We also add some new tools which we present in Sections 4.2 and 4.3. Finally, we discuss implementation and efficiency details in Section 4.4.

4.1 The DFT method

To obtain lower bounds on topological entropy, we compute a semi-conjugate symbolic dynamical system using the DFT method, which is based on the discrete Conley index. We use this method largely as a black box, and hence give only a rough sketch of the technique here and refer the reader to Day et al. [6] for the full details. At a high level, the method consists of three main steps:

- I. Discretize a compact subset of the phase space by constructing a grid \mathcal{G} of boxes, and compute a rigorous enclosure for the dynamics on these boxes as described in Section 2.3.
- II. Find a combinatorial index pair from these boxes and determine the Conley index for this pair using [19].
- III. From the index, compute a subshift to which the original system is semi-conjugate.

This approach is very general, and in principle could be applied to systems of arbitrary dimension. A major benefit to using this method here, however, is that in our setting it is completely automated. As long as one knows roughly where in the phase space the invariant objects of interest are, one can simply plug in the parameters and compute. See Section 4.4 for more details on the implementation.

An advantage to studying hyperbolic parameters of Hénon map (1) for $|b| \leq 1$ is that the nonwandering set is disconnected. This follows from Plykin theory, as discussed in [21, §7.9], from which we know that any connected trapping region (a region N with $f(N) \subseteq \text{int}(N)$) of the attractor has at least three holes. Considering a disc covering such a hole, we see that since the image each hole must strictly cover another hole, an iterate of this disk must eventually expand, which contradicts the area-preserving or area-shrinking of the maps we are considering. Thus, all such trapping regions must be disconnected.

Since the nonwandering set is disconnected, we can bypass much of the complication in the second step of the DFT method, that of finding an index pair. This is because the invariant set will be naturally isolated; at a fine enough resolution, the collection of boxes that cover the invariant set will already be separated into disjoint regions.

4.2 Reducing large subshifts

When applying the DFT method to complex systems, the resulting semi-conjugate subshift (step III above) is often very large. In fact, the main result in [6] is shown via a subshift with 247 symbols. While such large subshifts provide a topological entropy bound and information about the number of orbits of a given period, they often carry little *intuition* about the underlying structure of the dynamics. The question then becomes, how can we distill more useful information from these large subshifts to get a more intuitive understanding of the system?

The answer we propose here is to simplify the resulting subshift using amalgamations, as described in Section 3. Specifically, we add a final step to the DFT method, where we run the semi-conjugate subshift obtained through Algorithm 1. As we will see in Section 5, in practice this procedure can greatly reduce the number of symbols required to describe the system. In some cases this simplification reveals a simple underlying structure, such as connections between a handful of low-period orbits, which would otherwise be difficult to glean from the 200+ symbols of the original subshift. This simplification is also useful in comparing our results to previous work and conjectures; using Algorithm 1, one can attempt to amalgamate our computed subshift A to yield a target subshift B . We apply this technique in Section 5.1 to partially confirm a conjecture of Davis et al. [5]; see Figure 8 in particular.




Figure 3: Sensitivity of the entropy bound with respect to the aspect ratio (a) and the area (b) of the boxes. The map used was Hénon (1) with parameters (5.685974, -1), or plateau 14 from Section 5.1.

As a final note, recall that the subshift generated by the DFT method has a useful geometric interpretation: from [6] we know that we can associate a region N_i of the phase space to each symbol s_i of the subshift A , such that any trajectory $(\dots, s_{-1}, s_0, s_1, \dots)$ in A corresponds to a trajectory in the original system through the regions $(\dots, N_{-1}, N_0, N_1, \dots)$.⁶ Fortunately, amalgamation (and thus Algorithm 1) preserves this property in the following sense. If A' is derived from A by a sequence of amalgamations, each symbol s'_i of A' can be expressed as an amalgamation of symbols of A . Taking N'_i to be the union of the regions corresponding to these symbols, it is easy to see that the same trajectory property holds for A' .

4.3 Robustness and scaling

A natural concern for any approach which involves discretizing the phase space is the robustness of the method with respect to the choice of discretization. In our setting, this discretization is determined by the resolution (depth), aspect ratio,⁷ and the precise placement of the grid \mathcal{G} .

Ideally, slight changes in these scaling parameters would have at most minor effects on the resulting subshift or entropy bound, but there are several reasons why this is unreasonable to expect. Perhaps most obvious is that a subshift of finite type, and even the Conley index itself, is a discrete object, and one cannot expect any sort of continuity in the scaling parameters when the output itself is discrete. Moreover, the combinatorial isolation of the index pair depends on the topological properties of rigorous numerical bounds on the images of boxes, which can be very sensitive to the precise grid parameters. A simple example is in achieving disjoint regions; when trying to separate regions A and B

⁶In fact, these regions are connected components of the maximal invariant set of the index pair.

⁷We will use this term to refer to the shape of the boxes in higher dimensions as well.

of the phase space with $\text{dist}(A, B) = d$, using a grid whose boxes \mathcal{B}_i have width $y_k^i - x_k^i = 4d$ in each dimension k , great care must be taken in placing the grid, or the entire index pair could degenerate. More subtle is a situation where the combinatorial image of A does not overlap B , but does after a slight shift or rescaling of the grid. These issues raise the question of the practical robustness of the DFT method with respect to these scaling parameters.

To measure this robustness, we plot the computed entropy lower bound for a particular Hénon map against the area and the aspect ratio of the boxes in Figure 3. There we define $w_1 = y_1^i - x_1^i$ and $w_2 = y_2^i - x_2^i$ to be the width of a box \mathcal{B}_i in dimensions 1 and 2, respectively (recall that in our setting these are independent of i). The grid resolution may be defined as $-\frac{1}{2} \log_2 w_1 w_2$, or more generally as $-\frac{1}{n} \log_2 V(\mathcal{B}_i)$, where V denotes the n -dimensional volume. We choose this formula so that after normalizing by the volume of the initial box \mathcal{B} , the notions of resolution and depth align.

Of course, the DFT method may behave very differently on other maps, but the behavior shown is quite typical: the entropy lower bound is roughly monotone increasing with respect to the resolution, and (very) roughly unimodal with respect to the aspect ratio. This behavior is not surprising; the error from discretization decreases with the box area, and both extremes of the aspect ratio (i.e. $w_1 \gg w_2$ and $w_2 \gg w_1$) should result in essentially 1-dimensional information and hence a trivial entropy lower bound for most maps.

While the DFT method appears to be relatively robust with respect to the grid resolution, Figure 3(a) clearly shows a high sensitivity to the aspect ratio. Specifically, small changes in the aspect ratio at depth 9 resulted in large jumps in the entropy lower bound, even when close to the optimal ratio. While the behavior at depth 10 is somewhat more typical, this sensitivity is still something to keep in mind. In particular, for the sake of replication, care should be taken when altering and storing the aspect ratio.^{8,9}

As mentioned above, the DFT method gives an entropy bound which is roughly monotone increasing with the resolution. Similar plots shown in [6] suggested monotonicity, but had very few data points; it is likely that a continuous resolution scaling would fill in these plots to reveal a rough monotonicity in that case as well. This monotonicity is of course beneficial behavior, as one would like the precision of the bound to increase with the precision of the box covering, but it is especially useful in hyperbolic settings. By [21, Theorem 9.6.1], hyperbolic systems admit a finite Markov partition of the invariant set, and since the nonwandering set is disconnected (see Section 4.1), in theory this partition is obtainable when the boxes are small enough. Thus we expect to obtain the true entropy value at a high enough resolution, and by seeing where the entropy levels off we can be confident, though not certain, that our lower bound is the actual value. We will apply this intuition in Section 5.

⁸To find an appropriate aspect ratio for a given map, one can perform a scaling parameter exploration at a lower depth, similar to the one in Figure 3(a).

⁹Note also that one can scale a map in a non-constant fashion to better align it with the grid. Consider the map $f(x, y) = ((a - (\alpha^y x)^2 + by)\alpha^{-\alpha^y x}, \alpha^y x)$, which is just the Hénon map (1) conjugated by $g(x, y) = (\alpha^y x, y)$, and thus for large enough α would be more aptly analyzed by first conjugating by g^{-1} . More generally, such non-constant scaling may be used to focus on areas of the phase space where isolation is more difficult.

4.4 Implementation and efficiency

The computations in this paper were performed in Matlab, using Intlab [22], GAIO [8], and CHomP [4]. The machines used had memory between 1 and 2 gigabytes and clock speeds between 1 and 2.2 gigahertz. The runtimes varied, as we will describe below, but ranged from a few seconds at low grid resolutions to several hours for high resolutions.

As mentioned above, it is useful to observe how the entropy bound changes as we increase the grid resolution, but of course this procedure is not without cost. One would expect the runtime of the bound computation to increase, perhaps dramatically, with the resolution of the discretization. Empirically, the running time seems to grow very roughly as n^r where n is the dimension of the invariant set, and r is the resolution of the grid, as defined above.

Fortunately, as we have mentioned several times, the DFT method and the new techniques we add here are all completely automated. By this we mean that one needs only to specify the map and the parameters to be studied (and the discretization parameters discussed above), and the rest of the computation, all the way to the semi-conjugate subshift, is done without any further human action or input. Thus, although the computations may be time-consuming for high resolutions, it is computation time, not human time.

This lack of manual intervention enables vast explorations of map and discretization parameters. For this paper, a total of 58 parameter values of the Hénon map were studied: 43 in Section 5 and 15 in Section 5.1. For each of the 58 parameter values, we computed entropy bounds at 64 or 72 different resolutions, yielding a total of over 3900 separate computations. This exploration would have been infeasible without complete automation (or a large team of researchers).

5 Hyperbolic plateaus of Hénon

We now apply the methods outlined in Section 4 to the real-valued Hénon map

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

for parameter values (a, b) such that $f_{a,b}$ is (uniformly) hyperbolic. Note that this *excludes* the classical parameters $(1.4, 0.3)$; for rigorous topological lower bounds in the classical case, see [6] and [18].

Using the hyperbolic plateaus of Arai, we select representative parameter values to study for each plateau, as shown in Figure 4. For each plateau, we use the continuous resolution scaling approach mentioned in Section 4.3, to get a sense of how close our bounds might be to the actual values. The entropy bounds we compute constitute our main result, summarized in Theorem 5.1.




Figure 4: Hyperbolic plateaus for Hénon from [1], with a on the horizontal axis and b on the vertical axis. The label for each plateau is centered over the parameter values used to represent the plateau.

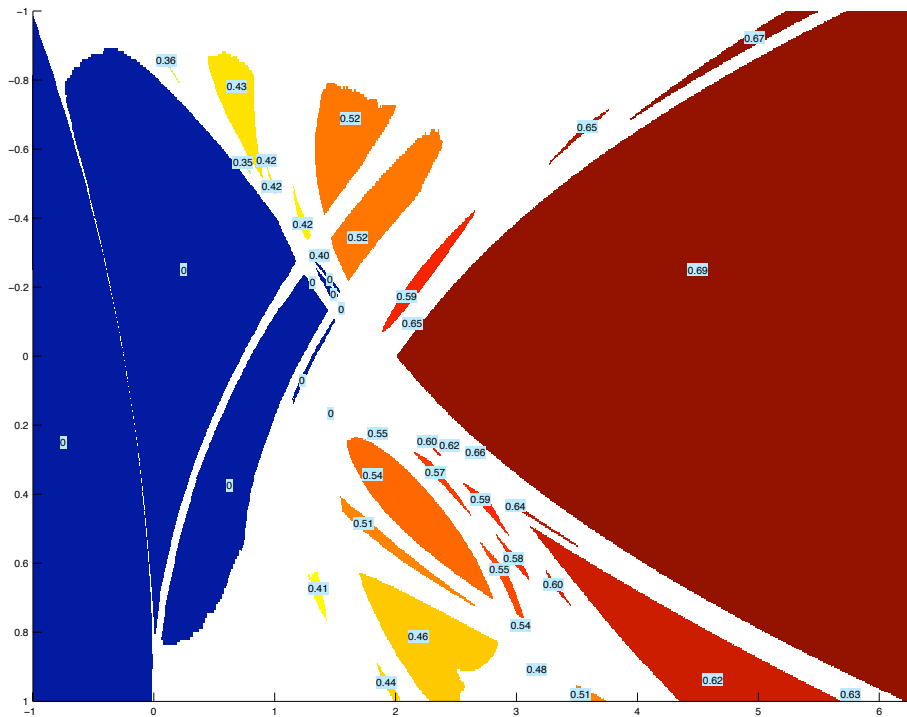


Figure 5: Rigorous lower bounds for topological entropy for the hyperbolic plateaus labeled 1 through 43 in Figure 4.

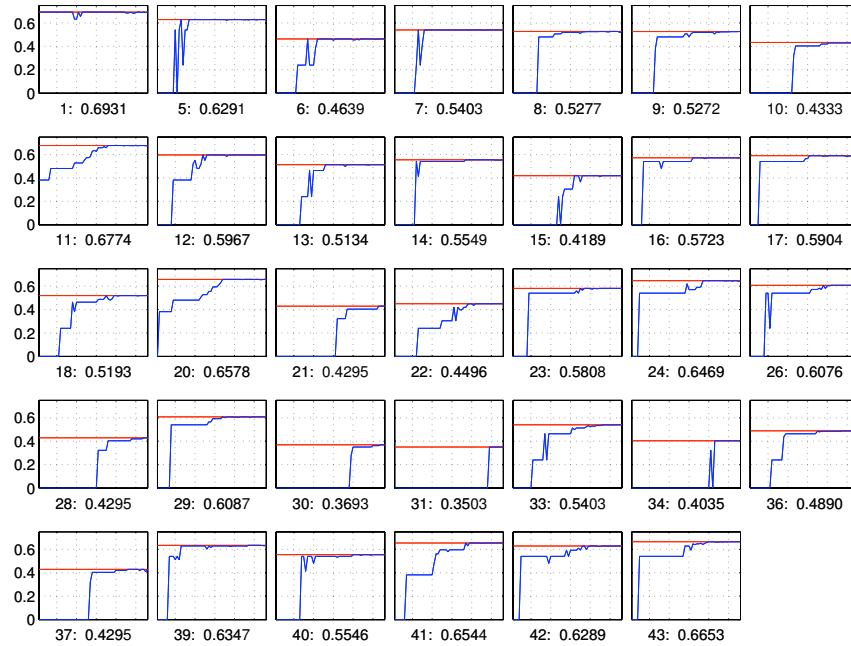


Figure 6: Entropy versus resolution. The red horizontal line in each represents the maximum lower bound computed (also printed below each plot). Plots are omitted for 0-entropy plateaus (2, 3, 4, 19, 25, 27, 32, 35, 38).

Theorem 5.1. Let $F_i = \{f_{a,b} \mid (a,b) \in R_i\}$, where R_i is the i th plateau in Figure 4. Then for all i and all $f \in F_i$ we have $h(f) \geq h_i$, where the h_i are defined below:

$$\begin{aligned}
 h_1 &= 0.6931 & h_5 &= 0.6291 & h_6 &= 0.4639 & h_7 &= 0.5403 & h_8 &= 0.5277 \\
 h_9 &= 0.5270 & h_{10} &= 0.4333 & h_{11} &= 0.6774 & h_{12} &= 0.5967 & h_{13} &= 0.5134 \\
 h_{14} &= 0.5549 & h_{15} &= 0.4189 & h_{16} &= 0.5723 & h_{17} &= 0.5904 & h_{18} &= 0.5193 \\
 h_{20} &= 0.6578 & h_{21} &= 0.4295 & h_{22} &= 0.4496 & h_{23} &= 0.5808 & h_{24} &= 0.6469 \\
 h_{26} &= 0.6076 & h_{28} &= 0.4295 & h_{29} &= 0.6087 & h_{30} &= 0.3693 & h_{31} &= 0.3503 \\
 h_{33} &= 0.5403 & h_{34} &= 0.4035 & h_{36} &= 0.4890 & h_{37} &= 0.4295 & h_{39} &= 0.6347 \\
 h_{40} &= 0.5546 & h_{41} &= 0.6544 & h_{42} &= 0.6289 & h_{43} &= 0.6653.
 \end{aligned}$$

Proof. For each R_i we selected $(a_i, b_i) \in R_i$ as a representative (these choices are shown in Figure 4). We then computed a combinatorial enclosure via interval arithmetic for f_{a_i, b_i} at different resolutions, and for each resolution we computed a rigorous lower bound for topological entropy using the DFT method (see Section 4.1). These bounds are summarized in Figure 6. Finally, by [1] we know each R_i is (uniformly) hyperbolic and so for each i we can apply the maximum lower bound achieved for (a_i, b_i) to all of R_i . \square

Figures 1 and 5 show an overview of our results. Note that the entropy values shown are merely lower bounds, and not necessarily the true values. Since we are computing these bounds for hyperbolic parameter values, however, we know from Section 4.3 that if our entropy lower bound levels off as the grid resolution increases, we have strong evidence that we have obtained the correct value. Typical index pairs from these computations are




Figure 7: Two index pairs from the computations. The black is the set P_0 while the colors correspond to the amalgamated regions after applying Algorithm 1.

shown in Figure 7, with colored regions corresponding to symbols in the resulting subshift after amalgamation using Algorithm 1.

In Figure 6, we show plots of entropy bounds computed versus the resolution, for each of the 43 parameter values with a nonzero bound. Using the above heuristic, it seems that for most of the plateaus, the bounds we computed should be exact or very close. A few notable exceptions are plateaus 9, 10, 21, 28, 30, and 31, since the plots do not seem to have leveled off, and we would expect many of these bounds to improve with further computations. For plateaus 33, 34, 36, 37, and 39, it is unclear whether the plot has leveled off. As we saw in Section 4.3, the DFT method is fairly robust with respect to the grid resolution, and the entropy lower bound is roughly monotonic in the resolution. The plots in Figure 6 reaffirm this, with very few exceptions.

While we have computed a large array of lower bounds, covering a vast portion of the parameter space of the Hénon map, a recent result of Arai gives a method of rigorously computing *exact* entropy values for (uniformly) hyperbolic Hénon [2]. In fact, in that work he computes values for plateaus numbered 5, 7, 11, and 12 in Figure 4, which match our lower bounds precisely. In that it computes exact entropy values, Arai’s method is certainly superior to the DFT method in the case we study here, and indeed it would be very interesting to use his method to test the accuracy of our other lower bounds. However, it is important to note that since Arai’s method in [2] relies heavily on hyperbolicity, it is not as generally applicable as the DFT method; in particular, it could not be directly applied to the Hénon map for the classical parameter values, studied in [6], as the map is not uniformly hyperbolic for those parameters.

While much of the parameter space in Figure 5 is covered by our lower bounds, there is still much of the parameter space which is not hyperbolic, or has not yet been proven to be hyperbolic. Thus, it remains in future work to lower-bound the entropy in the remaining white regions. At first this seems like a daunting or impossible task, since in the nonhyperbolic regions, we no longer have plateaus of topological entropy, and thus

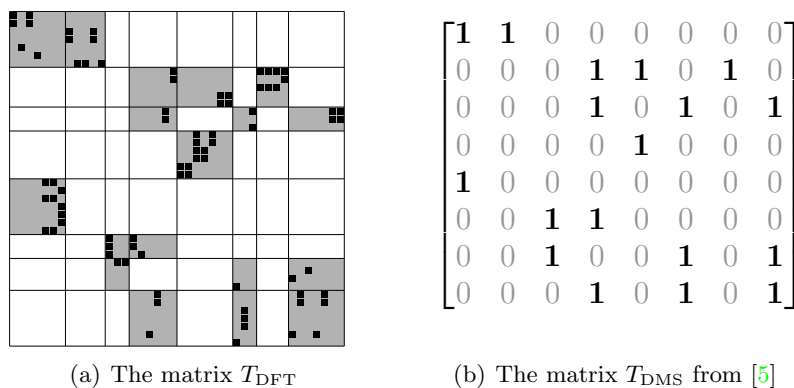


Figure 8: Amalgamation of the 42×42 symbol matrix obtained using the DFT method. The black squares in (a) denote the nonzero entries of T_{DFT} while the gray regions represent the amalgamated symbols, which one can easily see match T_{DMS} exactly.

cannot extend a bound from a single set of parameters to an entire region. Fortunately, the DFT method can be applied to intervals of parameter values $[a_1, a_2] \times [b_1, b_2]$, yielding a single lower bound which applied to the entire interval, as demonstrated in [6] and [9]. Thus, it would be of great interest to compute lower bounds on an interval tiling of the Hénon parameter space, and then compare these bounds to those computed here; as mentioned in Section 4.4, the automation of the DFT method would enable such parameter explorations.

5.1 Area-preserving Hénon maps

When $b = -1$, the Hénon maps are area-preserving and orientation-preserving. This case has been well studied, especially in the Physics literature. Starting in 1991, Davis, MacKay, and Sannami (DMS) [5] conjectured that Hénon was hyperbolic for three values of a (5.4, 5.59, 5.65) and conjectured conjugacies to symbolic dynamics for these three values as well. In 2002, de Carvalho and Hall in [7] replicated the results for $a = 5.4$ using a pruning approach. In 2004, Hagiwara and Shudo in [12] used a different pruning method to replicate all three of the values that DMS studied, and two more. They also give estimates of the topological entropy for $4 \leq a \leq 5.7$, which are displayed in Figure 10. Finally in 2007 some rigorous results appeared by Arai in [1], where he proved that there are 16 hyperbolic regions for $b = -1$, covering the parameters studied by DMS and the two others studied by Hagiwara and Shudo. Arai goes on in [2] to prove that the subshift conjectured by DMS for $a = 5.4$ is actually conjugate.

While our method cannot prove exact topological entropy values or conjugacies, we can attempt to verify that the entropy of the subshifts given in [5] are lower bounds, and perhaps show that the subshifts themselves are semi-conjugate. We focus first on the $a = 5.4$ case, where Davis et al. conjectured that f_a is conjugate to the subshift corresponding to the transition matrix T_{DMS} in Figure 8(b), which has topological entropy $h(T_{\text{DMS}}) \approx 0.6774$. Note that this is the same plateau as plateau 11 from the previous section.




Figure 9: Index pair for $a = 5.4$ and $b = -1$, with collections of regions labeled to indicate the symbols for the smaller, shift-equivalent symbol system.

Using our technique, we prove that $f_{5.4}$ is semi-conjugate to a 42×42 symbol matrix T_{DFT} , depicted in Figure 8(a). The topological entropy of this matrix is the same value, $h(T_{\text{DFT}}) \approx 0.6774$. The fact that $h(T_{\text{DMS}}) = h(T_{\text{DFT}})$ suggests that the subshifts given by T_{DMS} and T_{DFT} might be conjugate, and indeed we prove this in Theorem 5.2.

Theorem 5.2. *The map $f_{5.4}$ is semi-conjugate to the subshift T_{DMS} given in Figure 8(b). Moreover the symbols of T_{DMS} correspond to the regions labeled in Figure 9, which are the same regions conjectured by DMS.*

Proof. Applying Algorithm 1 to T_{DFT} , we obtain a strong shift equivalence between T_{DMS} and T_{DFT} , which shows that the corresponding subshifts are indeed conjugate. Moreover, the amalgamated vertices can be chosen so that we obtain the same partition that was used by Davis et al., which is shown in Figure 9, with the regions labeled so as to match the symbols (row indices) of T_{DMS} . \square

Our method also gives lower bounds which match the values conjectured by Davis et al. for $a = 5.59$ and $a = 5.65$, as well as the value conjectured by Hagiwara and Shudo for $a = 4.58$. In addition to these values, we also focus on 11 other values, which all together correspond to the first 15 area-preserving plateaus computed by Arai (the 16th is the maximal entropy plateau, which is plateau 1 in the previous section). Our results are summarized in the following theorem.




Figure 10: Topological entropy of the area-preserving Hénon maps, where $b = -1$. The estimates produced by Hagiwara and Shudo are in blue, and our rigorous lower bounds are in green.

Theorem 5.3. *The following entropy bounds hold for the Hénon maps $f_a = f_{a,-1}$. Here we write $h(f_{[a_0, a_1]}) \geq v$ to mean $\forall a \in [a_0, a_1], h(f_a) \geq v$.*

- | | |
|---|---|
| 1. $h(f_{[4.5383, 4.5386]}) \geq 0.6373$ | 2. $h(f_{[4.5388, 4.5430]}) \geq 0.6373$ |
| 3. $h(f_{[4.5624, 4.5931]}) \geq 0.6391$ | 4. $h(f_{[4.6189, 4.6458]}) \geq 0.6404$ |
| 5. $h(f_{[4.6694, 4.6881]}) \geq 0.6429$ | 6. $h(f_{[4.7682, 4.7993]}) \geq 0.6459$ |
| 7. $h(f_{[4.8530, 4.8604]}) \geq 0.6466$ | 8. $h(f_{[4.9666, 4.9692]}) \geq 0.6527$ |
| 9. $h(f_{[5.1470, 5.1497]}) \geq 0.6718$ | 10. $h(f_{[5.1904, 5.5366]}) \geq 0.6774$ |
| 11. $h(f_{[5.5659, 5.6078]}) \geq 0.6814$ | 12. $h(f_{[5.6343, 5.6769]}) \geq 0.6893$ |
| 13. $h(f_{[5.6821, 5.6858]}) \geq 0.6893$ | 14. $h(f_{[5.6859, 5.6860]}) \geq 0.6893$ |
| 15. $h(f_{[5.6917, 5.6952]}) \geq 0.6893$ | |

Proof. Using the DFT method (see Section 4.1), we computed rigorous lower bounds on topological entropy for a single a value for each plateau; the representatives chosen were the following: 4.5385, 4.5409, 4.5800, 4.6323, 4.6788, 4.7838, 4.8600, 4.9679, 5.1483, 5.4000, 5.5900, 5.6500, 5.6839, 5.6859, 5.6934. Combining these bounds with the hyperbolic plateaus computed in [1], and using Theorem 2.4, we can extend each bound to its corresponding plateau. \square

Figure 10 shows a plot of the lower bounds from Theorem 5.3, shown against the estimates computed by Shudo and Hagiwara in [12]. The 4 cases discussed above correspond to plateaus 3, 10, 11, and 12. For these plateaus, our lower bounds match the estimates exactly, and our bounds for plateaus 4, 5, and 6 are very close. This is roughly what one would expect given the resolution plots shown in Figure 11. An interesting trend we see in these data is that the algorithm performed better on the larger plateaus. This is perhaps because the stable and unstable manifolds seem to be more transverse the farther the parameters are from a bifurcation, making it easier to isolate the important regions of the phase space.

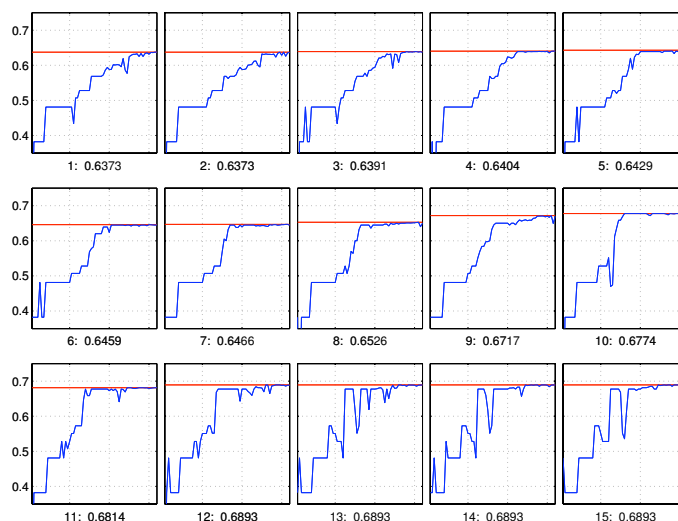


Figure 11: Entropy against resolution, as in Figure 6.

Acknowledgements

The author would like to thank John Smillie for all of his guidance and advice, Zin Arai for generously sharing his useful data, Sarah Day and Rodrigo Treviño for their helpful comments on numerous drafts of this paper, and Daniel Lepage for his help with the visualization in Figure 1.

References

- [1] Zin Arai. On hyperbolic plateaus of the Hénon map. *Experiment. Math.*, 16(2):181–188, 2007.
- [2] Zin Arai. On loops in the hyperbolic locus of the complex Hénon map and their monodromies. Preprint, 2008.
- [3] Rufus Bowen. Periodic points and measures for Axiom A diffeomorphisms. *Trans. Amer. Math. Soc.*, 154:377–397, 1971.
- [4] Computational Homology Project (CHomP), 2008. <http://chomp.rutgers.edu/>.
- [5] M. J. Davis, R. S. MacKay, and A. Sannami. Markov shifts in the Hénon family. *Phys. D*, 52(2-3):171–178, 1991.
- [6] Sarah Day, Rafael Frongillo, and Rodrigo Treviño. Algorithms for rigorous entropy bounds and symbolic dynamics. *SIAM Journal on Applied Dynamical Systems*, 7:1477–1506, 2008.
- [7] André de Carvalho and Toby Hall. How to prune a horseshoe. *Nonlinearity*, 15(3):R19–R68, 2002.
- [8] Michael Dellnitz, Gary Froyland, and Oliver Junge. The algorithms behind GAIO-set oriented numerical methods for dynamical systems. In *Ergodic theory, analysis, and efficient simulation of dynamical systems*, pages 145–174, 805–807. Springer, Berlin, 2001.
- [9] Rafael Frongillo and Rodrigo Treviño. Efficient automation of index pairs in computational conley index theory. *SIAM Journal on Applied Dynamical Systems*, 11:82, 2012.

- [10] Rafael M. Frongillo. On the hardness of state amalgamation. Preprint, 2014.
- [11] John Guckenheimer and Philip Holmes. *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*, volume 42 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 1990. Revised and corrected reprint of the 1983 original.
- [12] Ryouichi Hagiwara and Akira Shudo. An algorithm to prune the area-preserving Hénon map. *J. Phys. A*, 37(44):10521–10543, 2004.
- [13] Suzanne Lynch Hruska. A numerical method for constructing the hyperbolic structure of complex Hénon mappings. *Found. Comput. Math.*, 6(4):427–455, 2006.
- [14] Tomasz Kaczynski, Konstantin Mischaikow, and Marian Mrozek. *Computational homology*, volume 157 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 2004.
- [15] Douglas Lind and Brian Marcus. *An Introduction to Symbolic Dynamics and Coding*. Cambridge University Press, Cambridge, 1995.
- [16] M. Mazur and J. Tabor. Computational hyperbolicity. *Discrete and Continuous Dynamical Systems (DCDS-A)*, 29(3):1175–1189, 2011.
- [17] Konstantin Mischaikow and Marian Mrozek. Conley index. In *Handbook of Dynamical Systems*, pages 393–460. Elsevier, Berlin, 2002.
- [18] S. Newhouse, M. Berz, J. Grote, and K. Makino. On the estimation of topological entropy on surfaces. *Contemporary Mathematics*, 469:243–270, 2008.
- [19] Pawel Pilarczyk. *Homology Computation-Software and Examples*. Jagiellonian University, 1998. <http://www.pawelpilarczyk.com/homology.php>.
- [20] Joel W. Robbin and Dietmar Salamon. Dynamical systems, shape theory and the Conley index. *Ergodic Theory Dynam. Systems*, 8*(Charles Conley Memorial Issue):375–393, 1988.
- [21] Clark Robinson. *Dynamical systems*. Studies in Advanced Mathematics. CRC Press, Boca Raton, FL, 1995. Stability, symbolic dynamics, and chaos.
- [22] S.M. Rump. INTLAB - INTerval LABoratory. In Tibor Csendes, editor, *Developments in Reliable Computing*, pages 77–104. Kluwer Academic Publishers, Dordrecht, 1999. <http://www.ti3.tu-harburg.de/rump/>.
- [23] D. Sterling, H. R. Dullin, and J. D. Meiss. Homoclinic bifurcations for the Hénon map. *Phys. D*, 134(2):153–184, 1999.
- [24] R. F. Williams. Classification of subshifts of finite type. *Annals of Mathematics*, 98(1):120–153, 1973.