

Combinatorial-topological framework for the analysis of global dynamics

Cite as: Chaos 22, 047508 (2012); <https://doi.org/10.1063/1.4767672>

Submitted: 26 April 2012 . Accepted: 01 November 2012 . Published Online: 14 December 2012

Justin Bush, Marcio Gameiro, Shaun Harker, Hiroshi Kokubu, Konstantin Mischaikow, Ippei Obayashi, and Paweł Pilarczyk



View Online



Export Citation



CrossMark

ARTICLES YOU MAY BE INTERESTED IN

[Persistent topological features of dynamical systems](#)

Chaos: An Interdisciplinary Journal of Nonlinear Science **26**, 053105 (2016); <https://doi.org/10.1063/1.4949472>

[Persistent homology of time-dependent functional networks constructed from coupled time series](#)

Chaos: An Interdisciplinary Journal of Nonlinear Science **27**, 047410 (2017); <https://doi.org/10.1063/1.4978997>

[Homology and symmetry breaking in Rayleigh-Bénard convection: Experiments and simulations](#)

Physics of Fluids **19**, 117105 (2007); <https://doi.org/10.1063/1.2800365>

NEW!

Sign up for topic alerts
New articles delivered to your inbox

AIP
Publishing



Combinatorial-topological framework for the analysis of global dynamics

Justin Bush,^{1,a)} Marcio Gameiro,^{2,b)} Shaun Harker,^{1,c)} Hiroshi Kokubu,^{3,d)} Konstantin Mischaikow,^{1,4,e)} Ippei Obayashi,^{3,f)} and Paweł Pilarczyk^{5,g)}

¹*Department of Mathematics, Hill Center-Busch Campus, Rutgers, The State University of New Jersey, 110 Frelinghusen Rd, Piscataway, New Jersey 08854-8019, USA*

²*Instituto de Ciências Matemáticas e de Computação, Universidade de São Paulo, Caixa Postal 668, 13560-970 São Carlos, SP, Brazil*

³*Department of Mathematics/JST CREST, Kyoto University, Kyoto 606-8502, Japan*

⁴*BioMaPS, Hill Center-Busch Campus, Rutgers, The State University of New Jersey, 110 Frelinghusen Rd, Piscataway, New Jersey 08854-8019, USA*

⁵*Centre of Mathematics, University of Minho, Campus de Gualtar, 4710-057 Braga, Portugal*

(Received 26 April 2012; accepted 1 November 2012; published online 14 December 2012)

We discuss an algorithmic framework based on efficient graph algorithms and algebraic-topological computational tools. The framework is aimed at automatic computation of a database of global dynamics of a given m -parameter semidynamical system with discrete time on a bounded subset of the n -dimensional phase space. We introduce the mathematical background, which is based upon Conley's topological approach to dynamics, describe the algorithms for the analysis of the dynamics using rectangular grids both in phase space and parameter space, and show two sample applications.

© 2012 American Institute of Physics. [<http://dx.doi.org/10.1063/1.4767672>]

It is well established that multiparameter nonlinear dynamical systems exhibit extremely complex behavior. For many applications, especially multiscale problems or in settings in which precise measurements are difficult, an understanding of coarse but robust structures that exist over large ranges of parameter values is of greater importance than a detailed understanding of the fine structure. With this in mind, we discuss a new mathematical and computational framework for the analysis of the global dynamics of multiparameter nonlinear systems. Our approach is based on a finite combinatorial approximation of phase space, parameter space, and the nonlinear dynamics. This is used to obtain a description of the global dynamics in terms of acyclic directed graphs called Morse graphs. A rigorous understanding of the dynamics is obtained using the Conley index, an algebraic topological invariant. The resulting information is finite and presented in the form of graphs and algebraic invariants and thus can be easily queried. For this reason, we view our procedure as producing a database for the global nonlinear dynamics for a parameterized nonlinear system. We include a discussion concerning the computational complexity of our approach along with two simple illustrative examples.

bifurcation of invariant sets often forms the focal point for the qualitative study of these systems. However, the theoretical work of the last century makes clear that invariant sets can possess structure on all spatial and temporal scales and furthermore that these structures can vary dramatically over parameter sets, which can be as complicated as Cantor sets of positive measure.

These results need to be contrasted with available methods of analysis, the ability to make measurements, and the derivation of models. In the context of applications, the focus is often on understanding the dynamics of a particular parameterized family of nonlinear systems. Because of the nonlinearity, typically this analysis is heavily dependent on performing and interpreting numerical simulations. Furthermore, these results are often compared against experimental data which itself is limited to finite measurements that contain errors and are of limited precision. Finally, many mathematical models, especially those arising from multiscale systems, are heuristic in nature; that is, the nonlinearities are not derived from first principle, but rather through a series of approximations. This implies that the exact values produced by the model at particular parameter values cannot be expected to and are not intended to exactly match those of the physical system. In addition, there are many instances of models for which crucial parameter values are unknown with bounds that range over many orders of magnitude.

This gap between the potentially infinite complexity of invariant sets and the crude tools of analysis and measurement suggests that an alternative perspective in describing the global properties of multiparameter families of nonlinear dynamical systems may be of use. In this paper, we provide a review of our attempts to develop such a new perspective with a focus on the computational aspects of the approach.

To keep technicalities to a minimum, we consider a multiparameter dynamical system given in the form of a continuous map,

I. INTRODUCTION

Physical models of evolutionary processes are typically framed in terms of continuous state spaces, parameter spaces, and time. Understanding the existence, structure, and

^{a)}E-mail address: bushjc@math.rutgers.edu.

^{b)}E-mail address: gameiro@icmc.usp.br.

^{c)}E-mail address: sharker@math.rutgers.edu.

^{d)}E-mail address: kokubu@math.kyoto-u.ac.jp.

^{e)}E-mail address: mischaik@math.rutgers.edu.

^{f)}E-mail address: obayashi@math.kyoto-u.ac.jp.

^{g)}E-mail address: pawel.pilarczyk@math.uminho.pt.

$$f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n, \\ (x, z) \mapsto f_z(x) := f(x, z), \tag{1}$$

$$F : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n \times \mathbb{R}^m, \\ (x, z) \mapsto (f_z(x), z). \tag{2}$$

where \mathbb{R}^n is the phase space and \mathbb{R}^m is the parameter space. However, it should be noted that this is not a serious restriction. Techniques that are analogous to those described in this paper have been successfully employed to study the dynamics of ordinary differential equations,²¹ partial differential equations,⁶ infinite dimensional maps,⁵ fast-slow systems,⁹ and time series analysis.¹⁹ Let $X \subset \mathbb{R}^n$ be a compact subset of phase space that contains the dynamics of interest and let $Z \subset \mathbb{R}^m$ be a compact subset of parameter space, which contains the set of physically relevant parameters. Our goal is to provide a mathematically rigorous description of the global dynamics restricted to X for all parameter values in Z .

Recall that for a given parameter value $z \in Z$, $S_z \subset X$ is an *invariant set* under f_z if $f_z(S_z) = S_z$. Traditionally, invariant sets are the focal point for dynamical systems. In the approach we present here, they play a secondary role. Instead, we focus on *isolating neighborhoods*; these are compact sets $N \subset X$ such that

$$\text{Inv}(N, f_z) \subset \text{int}(N),$$

where $\text{Inv}(N, f_z)$ denotes the maximal invariant set contained in N and $\text{int}(N)$ denotes the interior of N . Simple arguments based on continuity show that if N is an isolating neighborhood for f_z , then it is an isolating neighborhood for $f_{z'}$ for all z' sufficiently close to z . An invariant set that is the maximal invariant set in an isolating neighborhood is called an *isolated invariant set*. Given an isolating neighborhood, information about the isolated invariant set can be extracted using the Conley index, which is discussed in greater detail in Secs. **IIC** and **IID**. For the moment, it is sufficient to make three remarks.

N1 One can associate a Conley index to any isolating neighborhood.

N2 If N and N' are isolating neighborhoods and $\text{Inv}(N, f_z) = \text{Inv}(N', f_z)$, then they have the same Conley index.

N3 If N is an isolating neighborhood for all z in a path connected subset of Z , then the Conley index associated with N is the same for all f_z .

For justification of these remarks and further information about the Conley index, see Ref. **20**. The theme of this work is that isolating neighborhoods are relatively easy to identify, their Conley indices can be computed, and thus we can obtain information about isolated invariant sets. **N2** suggests that this information is relatively insensitive to the numerical approximations used to identify the isolating neighborhood. Furthermore, **N3** implies that the dynamical information extracted using the Conley index is robust with respect to perturbations in parameter values.

II. MATHEMATICAL FRAMEWORK

Because we are interested in structures which are invariant with respect to perturbations in parameter space, it is convenient to consider the parameterized dynamical system,

Given $Z \subset \mathbb{R}^m$, we denote the restriction of F to $\mathbb{R}^n \times Z$ by $F_Z : \mathbb{R}^n \times Z \rightarrow \mathbb{R}^n \times Z$.

A. Approximating dynamics

We can only perform a finite number of calculations, thus we need to combinatorialize phase space, parameter space, and the map that generates the dynamics. The discretization of phase space and parameter space is done using the concept of a *grid*.²⁴ This consists of a finite collection \mathcal{X} of nonempty, compact subsets of X with the following properties:

- (a) $X = \cup_{\xi \in \mathcal{X}} \xi$;
- (b) $\xi = \text{cl}(\text{int}(\xi))$ for all $\xi \in \mathcal{X}$;
- (c) $\xi \cap \text{int}(\xi') = \emptyset$ for all $\xi \neq \xi'$.

The *diameter* of a grid \mathcal{X} is defined by

$$\text{diam}(\mathcal{X}) = \sup_{\xi \in \mathcal{X}} \text{diam}(\xi).$$

As shown in Ref. **15**, any compact metric space admits a grid of arbitrarily small diameter. For $\mathcal{A} \subset \mathcal{X}$, the set $\cup_{\xi \in \mathcal{A}} \xi \subset X$ is denoted by $|\mathcal{A}|$.

For the sake of simplicity, in this paper, we only consider grids with grid elements in the form of cubes or simplices. With this in mind, we restrict the regions of phase space $X \subset \mathbb{R}^n$ and parameter space $Z \subset \mathbb{R}^m$ to be sets that can be represented by cubical or simplicial grids \mathcal{X} and \mathcal{Z} , respectively.

To discretize the dynamics, we make use of a *combinatorial multivalued map* $\mathcal{F} : \mathcal{X} \rightrightarrows \mathcal{X}$ which assigns to each element of a grid $\xi \in \mathcal{X}$ a subset (possibly empty) $\mathcal{F}(\xi)$ of \mathcal{X} . With regard to the algorithms that are employed in the analysis of the dynamics, it is important to observe that a combinatorial multivalued map is equivalent to a finite directed graph with vertices \mathcal{X} and directed edges (ξ, ξ') whenever $\xi' \in \mathcal{F}(\xi)$. With this in mind, we will refer to \mathcal{F} as a multivalued map or a directed graph, whichever is more convenient or intuitive given the situation.

To understand the relationship between multivalued maps and nonlinear dynamics, consider a continuous function $f : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and a compact subset $X \subset \mathbb{R}^n$. Let \mathcal{X} be a grid for X . A combinatorial multivalued map $\mathcal{F} : \mathcal{X} \rightrightarrows \mathcal{X}$ is an *outer approximation* of f , if

$$f(\xi) \subset \text{int}(|\mathcal{F}(\xi)|) \quad \text{for all } \xi \in \mathcal{X}. \tag{3}$$

Given a grid \mathcal{X} , the minimal outer approximation of f is given by

$$\tilde{\mathcal{F}}(\xi) := \{\xi' \mid \xi' \cap f(\xi) \neq \emptyset\}$$

and any other outer approximation \mathcal{F} of f satisfies $\tilde{\mathcal{F}}(\xi) \subset \mathcal{F}(\xi)$ for all $\xi \in \mathcal{X}$.¹⁵

An important observation is the fact that if $\tilde{\mathcal{F}}$ is a minimal outer approximation of f , then there exists $\delta > 0$ such

that if $\|g(x) - f(x)\| < \delta$ for all $x \in X$ then $\tilde{\mathcal{F}}$ is an outer approximation of g . Another outer approximation \mathcal{F} for f , which is not minimal, will in general allow for a larger δ . In this sense, grids and outer approximations provide a robust approximation of dynamics (Fig. 1).

From the computational perspective, determining the minimal outer approximation is typically too expensive. In general, the best that can be done is to compute an approximation of $f(\xi)$ along with an error bound ε which may or may not be small. With this information, one can construct an outer approximation that satisfies the following condition:

$$\{\xi' \in \mathcal{X} \mid \xi' \cap B_\varepsilon(f(\xi)) \neq \emptyset\} \subset \mathcal{F}(\xi).$$

The focus of this paper is on parameterized dynamical systems (2) for which we have chosen grids \mathcal{X} and \mathcal{Z} for X and Z , the regions of interest in phase space and parameter space, respectively. For each $\zeta \in \mathcal{Z}$, let $\mathcal{F}_\zeta : \mathcal{X} \rightrightarrows \mathcal{X}$ be an outer approximation of $F_\zeta : \mathbb{R}^n \times \zeta \rightarrow \mathbb{R}^n \times \zeta$ restricted to the grid $\mathcal{X} \times \zeta := \{\xi \times \zeta \mid \xi \in \mathcal{X}\}$. By definition, this implies that

$$f(\xi, \zeta) \subset \text{int}(|\mathcal{F}_\zeta(\xi)|) \quad \text{for all } \xi \in \mathcal{X}.$$

To understand how \mathcal{F}_ζ acts as an approximation of the underlying dynamics induced by Eq. (1), we state the following proposition which follows directly from the definition of an outer approximation.

Proposition 2.1. Let \mathcal{F}_ζ be an outer approximation for F_ζ . Consider any $x \in \xi$ and any ordered sequence of parameter values (z_1, z_2, \dots, z_T) where $\{z_i \in \zeta \mid i = 1, \dots, T\}$. Define

$$x_{i+1} := f_{z_i}(x_i),$$

where $x_0 = x$ and choose $\xi_i \in \mathcal{X}$ such that $x_i \in \xi_i$. Then the ordered sequence $(\xi, \xi_1, \dots, \xi_{T+1})$ is a path in the directed graph \mathcal{F}_ζ .

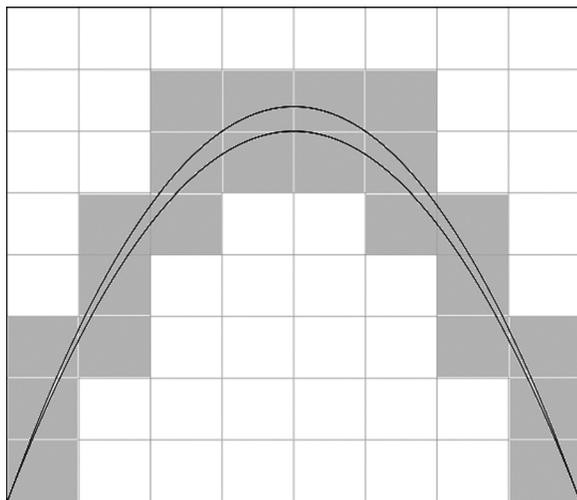


FIG. 1. An illustration of a minimal outer approximation. Depicted is the logistic map, $f(x) = rx(1 - x)$, with the region between the parabolas corresponding to $r \in [3, 3.2]$. For each grid element in the domain (in this case one of the eight intervals on the x -axis), the outer approximation of the grid element under f for the chosen values of r is the set of intervals in the range indicated by the shaded boxes above the domain grid element. This outer approximation is minimal because the set of shaded boxes is the fewest required to cover the graph of f for all values of $r \in [3, 3.2]$.

B. Extracting nonrecurrent dynamics

Assume that the grids \mathcal{X}, \mathcal{Z} are chosen and for a fixed $\zeta \in \mathcal{Z}$ an outer approximation $\mathcal{F}_\zeta : \mathcal{X} \rightrightarrows \mathcal{X}$ has been computed. The first step in using our approximation scheme to understand the dynamics generated by Eq. (1) is to identify the nonrecurrent dynamics.

Given the directed graph \mathcal{F}_ζ and $\mathcal{N} \subset \mathcal{X}$, the associated subgraph is the directed graph $\mathcal{F}_\zeta|_{\mathcal{N}} : \mathcal{N} \rightrightarrows \mathcal{N}$ consisting of the vertices $\{\xi \in \mathcal{N}\}$ and edges $\{(\xi, \xi') \mid \xi, \xi' \in \mathcal{N}, \xi' \in \mathcal{F}_\zeta(\xi)\}$. A directed graph is *invariant* if each vertex is both the head of at least one edge and the tail of at least one edge.

Proposition 2.2. Let $\mathcal{S}_\zeta \subset \mathcal{X}$ be the maximal invariant subgraph of \mathcal{F}_ζ . Then,

$$\text{Inv}(X, f_z) \subset |\mathcal{S}_\zeta|,$$

for all $z \in \zeta$.

The proof follows directly from Proposition 2.1. A consequence of Proposition 2.2 is that we have identified the portion of phase space on which the asymptotic dynamics takes place for all parameter values $z \in \zeta$. The next step is to identify the relevant dynamical structures which, likewise, are invariant for all $z \in \zeta$.

Given the directed graph \mathcal{F}_ζ , two elements $\xi, \xi' \in \mathcal{X}$ belong to the same *strongly connected path component* if there exist nontrivial paths from ξ to ξ' and ξ' to ξ .

Definition 2.3. Given a directed graph \mathcal{F}_ζ , let P_ζ be an index set for the collection of all strongly connected path components. The set

$$\{\mathcal{M}_\zeta(p) \subset \mathcal{X} \mid p \in P_\zeta\}$$

is the *Morse decomposition* of \mathcal{F}_ζ . The individual strongly connected path components are called *Morse sets*.

The following proposition summarizes basic facts about Morse sets, Morse decompositions, and an additional order structure on P_ζ . The proof follows from direct applications of the definition of strongly connected path components, the definition of maximal invariant subgraph, and Proposition 2.1.

Proposition 2.4. Consider an outer approximation $\mathcal{F}_\zeta : \mathcal{X} \rightrightarrows \mathcal{X}$ of F_ζ with maximal invariant subgraph \mathcal{S}_ζ . Then,

1. $\mathcal{M}_\zeta(p) \subset \mathcal{S}_\zeta$ for all $p \in P_\zeta$.
2. If $p \neq q$, then $\mathcal{M}_\zeta(p) \cap \mathcal{M}_\zeta(q) = \emptyset$
3. If $\xi \in \mathcal{S}_\zeta \setminus \cup_{p \in P_\zeta} \mathcal{M}_\zeta(p)$, then there exist $p, q \in P_\zeta$ and a path in \mathcal{F}_ζ that begins in $\mathcal{M}(p)$, passes through ξ , and ends in $\mathcal{M}(q)$. Note: If such a path exists for some $p, q \in P_\zeta$ then we write $q <_\zeta p$.
4. If $x \in \xi$ and $\xi \in \mathcal{S}_\zeta \setminus \cup_{p \in P_\zeta} \mathcal{M}_\zeta(p)$, then for all $z \in \zeta$ the point x is not a recurrent point of f_z restricted to X .
5. Under the relationship \leq_ζ , defined above, P_ζ is a partially ordered set.

Proposition 2.5 (See Theorem 4.1 in Ref. 15). Let $\{\mathcal{M}_\zeta(p) \subset \mathcal{X} \mid p \in P_\zeta\}$ be the Morse decomposition for the outer approximation $\mathcal{F}_\zeta : \mathcal{X} \rightrightarrows \mathcal{X}$ of F_ζ . Then for all $p \in P_\zeta$, $|\mathcal{M}_\zeta(p)|$ is an isolating neighborhood for f_z for all $z \in \zeta$.

Recall that given a partially ordered set (P, \leq) , we say that q covers p if from the relation $q \leq r \leq p$ it follows that either $q = r$ or $r = p$.

Definition 2.6. The Morse graph MG_ζ of \mathcal{F}_ζ is the acyclic directed graph with nodes consisting of the Morse sets and directed edges $\mathcal{M}_\zeta(p) \rightarrow \mathcal{M}_\zeta(q)$ if and only if q covers p in $(P_\zeta, \leq \zeta)$.

C. Identifying recurrent dynamics

For each grid element ζ in parameter space, the associated Morse graph MG_ζ provides rigorous information about the nonrecurrent dynamics and potential information about the recurrent dynamics that is valid over all parameter values in $|\zeta| \subset Z$. In particular, if recurrent dynamics occurs for some parameter value then it must occur within a region determined by a Morse set. We now describe the Conley index, which is an algebraic topological tool that can provide information about the recurrent dynamics.

We begin our description with some elementary topological constructions. Recall that a *pointed topological space* is a pair (V, v_0) where V is a topological space and $v_0 \in V$ is a distinguished point, sometimes called the *basepoint*. A continuous map $g : (V, v_0) \rightarrow (V, v_0)$ on a pointed topological space is a continuous map from V to V with the property that $g(v_0) = v_0$, i.e., the base point is a fixed point of the map.

Consider an arbitrary continuous map $g : Y \rightarrow Y$ defined on a locally compact metric space and a pair of compact subsets of Y denoted by $P = (P_1, P_0)$ with $P_0 \subset P_1$. Let $(P_1/P_0, [P_0])$ denote the pointed topological space where P_1/P_0 is the quotient space²⁵ obtained by collapsing P_0 to a single point denoted by $[P_0]$. Define $g_P : (P_1/P_0, [P_0]) \rightarrow (P_1/P_0, [P_0])$ by

$$g_P([x]) = \begin{cases} g(x) & \text{if } x, g(x) \in P_1 \setminus P_0 \\ [P_0] & \text{otherwise} \end{cases}.$$

Definition 2.7. A pair of compact sets $P = (P_1, P_0)$ is an *index pair* for g if

1. g_P is continuous, and
2. $\text{cl}(P_1 \setminus P_0)$ is an isolating neighborhood under g .

The induced map g_P is called an *index map*.

This definition of index pair and index map is due to Robbin and Salamon.²⁶ Using their characterization of index pairs and the nice topology of grids the proof of the following result, which implies that we can find index pairs associated with each of the regions that contain the potentially recurrent dynamics, is fairly straightforward.

Proposition 2.8. Let $\{\mathcal{M}_\zeta(p) \subset \mathcal{X} | p \in P_\zeta\}$ be the Morse decomposition for the outer approximation $\mathcal{F}_\zeta : \mathcal{X} \rightrightarrows \mathcal{X}$ of F_ζ . Assume $|\mathcal{S}_\zeta| \subset \text{int}(X)$. Let $P = (P_1, P_0)$ be defined by

$$P_1 := \left| \mathcal{F}_\zeta(\mathcal{M}_\zeta(p)) \right| \quad \text{and} \quad P_0 := \left| \mathcal{F}_\zeta(\mathcal{M}_\zeta(p)) \setminus \mathcal{M}_\zeta(p) \right|.$$

Then for all $z \in \zeta$, P is an index pair for f_z .

For the benefit of the reader who is unfamiliar with algebraic topology, we provide an informal description of the

purpose of the construction of the index map before completing our discussion of the Conley index. Assume that $g_P : (P_1/P_0, [P_0]) \rightarrow (P_1/P_0, [P_0])$ is an index map for the index pair $P = (P_1, P_0)$. Let $\dot{P} := P_1 \setminus P_0$. Observe that if $x \in \dot{P}$ and $g(x) \in \dot{P}$, then restricted to \dot{P} the maps g and g_P are the same. In particular, since by definition $\text{cl}(\dot{P})$ is an isolating neighborhood under g , $\text{Inv}(\text{cl}(\dot{P}), g) \subset \text{int}(\dot{P})$ and hence $\text{Inv}(\dot{P}, g)$ is equivalent to $\text{Inv}(\dot{P}, g_P)$. This is the motivation for turning our attention to studying the dynamics of g_P .

The most basic question one can ask is whether $\text{Inv}(\text{cl}(\dot{P}), g)$ is nonempty. Let us relate this to $g_P : (P_1/P_0, [P_0]) \rightarrow (P_1/P_0, [P_0])$. Since this is a continuous map on a pointed topological space, $g_P([P_0]) = [P_0]$. Assume $\text{Inv}(\text{cl}(\dot{P}), g) = \emptyset$. Then for any $x \in \dot{P}$, there exists a positive integer n such that $g^n(x) \notin \dot{P}$, which implies that $g_P^n(x) = [P_0]$. Compactness of $\text{cl}(\dot{P})$ then implies the existence of \bar{n} such that $g_P^{\bar{n}}(P_1/P_0) = [P_0]$. In words, after a sufficient number of iterates the dynamics of g_P on the quotient space $(P_1/P_0, [P_0])$ is a constant map. Observe that this argument suggests (though does not prove) that we can reduce the question of the existence of a nontrivial invariant set to showing that no iterate of g_P is the constant map. Algebraic topology, in particular homology, is used to demonstrate this last point.

In this context, homology^{12,14} is an assignment to any pointed topological space, in particular, an index pair $(P_1/P_0, [P_0])$, an infinite sequence of abelian groups, $H_k(P_1/P_0, [P_0])$, $k = 0, 1, 2, \dots$, and to a continuous map on that space, in particular, the associated index map $g_P : (P_1/P_0, [P_0]) \rightarrow (P_1/P_0, [P_0])$, an infinite sequence of group endomorphisms,

$$g_{P,k} : H_k(P_1/P_0, [P_0]) \rightarrow H_k(P_1/P_0, [P_0]), \quad k = 0, 1, 2, \dots$$

This sequence of group endomorphisms is a representative of the Conley index (where the equivalence relation is shift equivalence, defined later) for the isolating neighborhood $\text{cl}(\dot{P})$ under g .

While this may appear rather formidable, a variety of algebraic observations and constraints can be used to simplify this discussion. First, because the grids being considered are made up of cubes or simplices, only finitely many of the homology groups are nontrivial and the nontrivial homology groups are all finitely generated. Furthermore, one can compute homology groups using finite fields, e.g. \mathbb{Z}_2 or more generally \mathbb{Z}_p where p is a prime number, or rational numbers \mathbb{Q} . In this case, the homology groups are finite dimensional vector spaces and the linear maps $g_{P,k}$ can be represented as matrices. If the computations are done using \mathbb{Q} , then the nonzero eigenvalues $\bar{\sigma}_k$ of $g_{P,k}$ can be used as a representative of the Conley index.

An important observation is that given an index pair as in Proposition 2.8, the induced map on homology of an associated index map can be computed using \mathcal{F}_ζ .^{11,14,22} In summary, we have used the finite multivalued approximation of the map f to identify regions in phase space within which the existence of recurrent dynamics is suggested and we have transformed the problem of rigorously identifying the existence and structure of this recurrent dynamics to an algebraic problem. A more formal existence statement and the

most fundamental result associated with the Conley index is the following (see Ref. 20 for further discussion and references).

Theorem 2.9. *Let $g_{P,k} : H_k(P_1/P_0, [P_0]) \rightarrow H_k(P_1/P_0, [P_0])$ be induced by an index map. If for some $k \in 0, 1, 2, \dots$, $g_{P,k}$ is not nilpotent, then*

$$\text{Inv}(\text{cl}(P_1 \setminus P_0), g) \neq \emptyset.$$

The power of Theorem 2.9 is that it allows one to establish the existence of nonempty invariant sets by checking a computable algebraic invariant.

Though we make no attempt to indicate how or why, we note that in addition to guaranteeing existence homology can also be used to provide lower bounds on the structure of the invariant sets. The following theorem provides a simple example of how the Conley index can be used to extract more detailed information about the dynamics that is robust with respect to perturbations in parameter space.

Theorem 2.10. *Let $\mathcal{M}_\zeta(p)$ be a Morse set for the outer approximation $\mathcal{F}_\zeta : \mathcal{X} \rightrightarrows \mathcal{X}$ of F_ζ . Then,*

$$\bar{\sigma}_0 = \emptyset \quad \text{or} \quad \bar{\sigma}_0 = \left\{ e^{2\pi i k/T} \mid k = 0, \dots, T-1 \right\} \text{ for some } T > 0.$$

In the latter case,

$$|\mathcal{M}_\zeta(p)| = \bigcup_{i=0}^{T-1} N_i,$$

where $\{N_i \mid i = 0, \dots, T-1\}$ are mutually disjoint compact sets with the property that

$$F_\zeta(N_i) \subset N_{i+1}, \quad i = 0, \dots, T-1$$

and

$$F_\zeta(N_{T-1}) \subset N_0.$$

In particular, given any sequence of parameter values $\{z_j \mid j = 0, 1, 2, \dots\} \subset \mathcal{Z}$, any $x_0 \in N_0$ and $x_{j+1} := f_{z_j}(x_j)$, we have

$$x_{j+1} \in N_k, \quad \text{where } k = j + 1 \pmod T.$$

The proof of this theorem follows from [Ref. 2, Proposition 5.8] and Proposition 2.1.

Definition 2.11. The Conley-Morse graph CMG_ζ of \mathcal{F}_ζ consists of the Morse graph MG_ζ of \mathcal{F}_ζ along with the Conley index associated with each Morse set $\mathcal{M}_\zeta(p), p \in P_\zeta$.

D. Classifying dynamics over parameter space

The discussion in Secs. II A and II C is restricted to the dynamics of F_ζ , where $\zeta \in \mathcal{Z}$ is a single grid element in parameter space. Since the results are valid for every $\zeta \in \mathcal{Z}$, this provides a rigorous description of the dynamics for every point $z \in \mathcal{Z}$. What remains to be discussed is how the dynamics over different grid points $\zeta, \zeta' \in \mathcal{Z}$ are related. We begin by defining a relationship between the Morse sets.

Definition 2.12. Let $\zeta, \zeta' \in \mathcal{Z}$ such that $\zeta \cap \zeta' \neq \emptyset$. The clutching graph $\mathcal{I}(\zeta, \zeta')$ is the bipartite graph with vertices $P_\zeta \cup P_{\zeta'}$ and with edges

$$(p, q) \in P_\zeta \times P_{\zeta'} \quad \text{if and only if} \quad \mathcal{M}_\zeta(p) \cap \mathcal{M}_{\zeta'}(q) \neq \emptyset.$$

Proposition 2.13. Assume there is a unique edge (p, q) in the clutching graph $\mathcal{I}(\zeta, \zeta')$ that has either p or q as its endpoint. Then the Conley index of $|\mathcal{M}_\zeta(p)|$ under F_ζ is the same as the Conley index of $|\mathcal{M}_{\zeta'}(q)|$ under $F_{\zeta'}$.

Proof. Let $z \in \zeta \cap \zeta'$. By Proposition 2.5, $|\mathcal{M}_\zeta(p)|$ and $|\mathcal{M}_{\zeta'}(q)|$ are isolating neighborhoods. Let $S_z := \text{Inv}(|\mathcal{M}_\zeta(p)|, f_z)$ and $S'_z := \text{Inv}(|\mathcal{M}_{\zeta'}(q)|, f_z)$. Observe that it is sufficient to show that $S_z = S'_z$, since the result then follows from **N2**.

With this in mind, suppose $S_z \neq S'_z$. Without loss of generality, we can assume that there exists $x \in S_z \setminus S'_z$. This implies that there is a grid element $\xi \in \mathcal{X}$ such that $x \in \xi \in \mathcal{M}_\zeta(p) \setminus \mathcal{M}_{\zeta'}(q)$. Since $x \in S_z$, ξ belongs to a strongly connected path component and hence belongs to $\mathcal{M}_{\zeta'}(r)$ for some $r \in P_{\zeta'}$ where $r \neq q$. This implies that the clutching graph $\mathcal{I}(\zeta, \zeta')$ contains the edge (p, r) , contradicting the uniqueness of the edges with endpoint p . \square

The key step in the proof of Proposition 2.13 is **N2**. Observe that the validity of **N2** is not obvious. In general, the index pairs $P = (P_1, P_0)$ defined by

$$P_1 := \left| \mathcal{F}_\zeta \left(\mathcal{M}_\zeta(p) \right) \right| \quad \text{and} \quad P_0 := \left| \mathcal{F}_\zeta \left(\mathcal{M}_\zeta(p) \right) \setminus \mathcal{M}_\zeta(p) \right|$$

and $P' = (P'_1, P'_0)$ defined by

$$P'_1 := \left| \mathcal{F}_{\zeta'} \left(\mathcal{M}_{\zeta'}(q) \right) \right| \quad \text{and} \quad P'_0 := \left| \mathcal{F}_{\zeta'} \left(\mathcal{M}_{\zeta'}(q) \right) \setminus \mathcal{M}_{\zeta'}(q) \right|$$

will be different and hence the induced index maps,

$$f_{P,k} : H_k(P_1/P_0, [P_0]) \rightarrow H_k(P_1/P_0, [P_0])$$

and

$$f_{P',k} : H_k(P'_1/P'_0, [P'_0]) \rightarrow H_k(P'_1/P'_0, [P'_0])$$

will provide different representations of the Conley index. Thus to explain **N2** requires a discussion of the equivalence classes used to define the Conley index. This is best done in a fairly general setting, so consider two functions,

$$a : V \rightarrow V \quad \text{and} \quad b : W \rightarrow W,$$

where V and W are both either finitely generated abelian groups or finite dimensional vector spaces, and a and b are group homomorphisms or linear maps.

Definition 2.14. The maps a and b are shift equivalent if there exist morphisms,

$$r : V \rightarrow W \quad \text{and} \quad s : W \rightarrow V$$

such that

$$b \circ r = r \circ a \quad \text{and} \quad s \circ b = a \circ s$$

and a positive integer n such that

$$s \circ r = a^n \quad \text{and} \quad r \circ s = b^n.$$

As is shown in Ref. 8, it follows that if $f_{P,k}$ and $f_{P',k}$ are constructed as above, then $f_{P,k}$ and $f_{P',k}$ are shift equivalent. More generally, when we indicate that two Conley indices agree, then we mean that the representative index maps are shift equivalent.

Proposition 2.13 motivates the following definition.

Definition 2.15. Fix grids \mathcal{X} and \mathcal{Z} and outer approximations $\mathcal{F}_\zeta : \mathcal{X} \rightrightarrows \mathcal{X}$ for all $\zeta \in \mathcal{Z}$. Two Morse sets $M_\zeta(p)$ and $M_{\zeta'}(p')$ belong to the same *Morse continuation class*, if there exists a sequence of grid elements $\{\zeta_i | i = 0, \dots, I\}$ with $\zeta_0 = \zeta$, $\zeta_I = \zeta'$ and indexing elements $\{p_i \in P_{\zeta_i} | i = 0, \dots, I\}$ with $p_0 = p$, $p_I = p'$ such that for all $i = 0, \dots, I - 1$, there exists a unique edge (p_i, p_{i+1}) in the clutching graph $\mathcal{I}(\zeta_i, \zeta_{i+1})$ that has either p_i or p_{i+1} as its endpoint.

N3 combined with Proposition 2.13 leads to the following result.

Corollary 2.16. Let $M_\zeta(p)$ and $M_{\zeta'}(p')$ belong to the same continuation class. Let $z \in \zeta$ and $z' \in \zeta'$. Then the Conley index of $|M_\zeta(p)|$ under f_z is the same as the Conley index of $|M_{\zeta'}(p')|$ under $f_{z'}$.

Remark 2.17. It is possible that $M_\zeta(p)$ and $M_\zeta(q)$, $q, p \in P_\zeta$, belong to the same continuation class even if $p \neq q$. For example, consider a bistable system in the plane with one isolated stable equilibrium at $(1, 0)$ and the other at $(-1, 0)$. If we take $\theta \in S^1$ to be a parameter that merely rotates the system by $\theta/2$, then for sufficiently fine grids in parameter and phase space the two equilibria will lie in the same Morse continuation class. The intuitive picture is that by starting at $(1, 0)$ and following the equilibrium as θ moves around S^1 one ends up at $(-1, 0)$. Insofar as the discretization accurately reflects this continuous picture of the dynamics, then there will be continuation between the Morse sets containing each of the equilibria.

Extending the idea of continuation classes to Morse graphs is slightly more subtle.² Assume that the clutching graph $\mathcal{I}(\zeta, \zeta')$ has the property that each node is the endpoint of exactly one edge. This defines a bijection,

$$\begin{aligned} b_{\zeta, \zeta'} : P_\zeta &\rightarrow P_{\zeta'} \\ p &\mapsto q \end{aligned} \tag{4}$$

if (p, q) is an edge of $\mathcal{I}(\zeta, \zeta')$.

Definition 2.18. Fix grids \mathcal{X} and \mathcal{Z} and outer approximations $\mathcal{F}_\zeta : \mathcal{X} \rightrightarrows \mathcal{X}$ for all $\zeta \in \mathcal{Z}$. Two Conley-Morse graphs CMG_ζ and $CMG_{\zeta'}$ belong to the same *Conley-Morse graph continuation class* if there exists a sequence of grid elements $\{\zeta_i | i = 0, \dots, I\}$ with $\zeta_0 = \zeta$, $\zeta_I = \zeta'$ and indexing elements $\{p_i \in P_{\zeta_i} | i = 0, \dots, I\}$ with $p_0 = p$, $p_I = p'$ such that for all $i = 0, \dots, I - 1$ the bijection,

$$b_{\zeta_i, \zeta_{i+1}} : (P_{\zeta_i}, \leq_{\zeta_i}) \rightarrow (P_{\zeta_{i+1}}, \leq_{\zeta_{i+1}})$$

is a directed graph isomorphism.

Remark 2.19. To each Conley-Morse graph continuation class, we associate three distinct types of information:

- The Morse graph, which provides information about the structure of the non-recurrent dynamics. More precisely, the partial order obtained from the Morse graph by taking the

transitive closure constrains the non-recurrent dynamics. In particular, for every trajectory with α -limit set and ω -limit set lying in distinct Morse sets, the Morse set containing the α -limit set must be greater than the Morse set containing the ω -limit set in the partial order. In this way, the Morse graph can be understood as giving a schematic picture of the dynamics in phase space away from any recurrent dynamics.

- The Conley indices of the Morse sets, which provide information about the structure of the recurrent dynamics.
- The set of parameter grid elements whose Conley-Morse graphs belong to the Conley-Morse continuation class. One can understand continuation classes as identifying the region in parameter space where the identified recurrent and non-recurrent dynamics occurs at the scale of the computation. In particular, this provides a lower bound on the set of parameter values at which the dynamics must occur. We use the number of parameter grid elements to measure the *size* of the continuation class.

Remark 2.20. The information about the dynamics provided by our approach can be viewed as a database of dynamics for the multiparameter nonlinear dynamical system (1) restricted to the region of phase space $X \subset \mathbb{R}^n$ and parameter space $Z \subset \mathbb{R}^m$, where the minimal levels of resolution are determined by the diameters of the grids \mathcal{X} and \mathcal{Z} . In particular, we can construct the *continuation graph*; that is, a graph whose nodes consist of the Conley-Morse graph continuation classes and whose edges consist of the clutching graph information between representative Conley-Morse graphs. This type of information is exhibited in Figure 2 for the overcompensatory Leslie model,³⁰

$$\begin{bmatrix} x \\ y \end{bmatrix} \mapsto \begin{bmatrix} x(\theta_1 x + \theta_2 y)e^{-0.1(x+y)} \\ 0.7x \end{bmatrix}. \tag{5}$$

To limit the information to a comprehensible amount, the upper left corner indicates the continuation graph associated with the 26 largest Conley-Morse graph continuation classes.

E. Grid refinements, bifurcations, and structural stability

The extent to which computations can resolve dynamics at fine scales is affected by a variety of factors, including the method for construction of combinatorial multivalued maps, and the geometry and sizes of the grids in phase space and parameter space. In general, at finer grids some Morse sets may be split and the finer structure of dynamics in their regions may be revealed. As a consequence, the complexity of the Conley-Morse graphs depends on the grid size. Namely, if the parameter value is fixed, the Conley-Morse graph on a finer grid may be larger than that on a coarser grid, as long as Morse sets with non-trivial Conley index are concerned. To be more precise, recall that the Conley-Morse graph and its subgraph consisting of Morse sets with non-trivial Conley index are posets (partially-ordered sets). The poset of the subgraph on a finer grid is projected *onto* the poset on a coarser grid by a map naturally induced by the grid refinement.

The same relation holds true for parameter grid refinement. Namely, for a parameter grid element and one of its

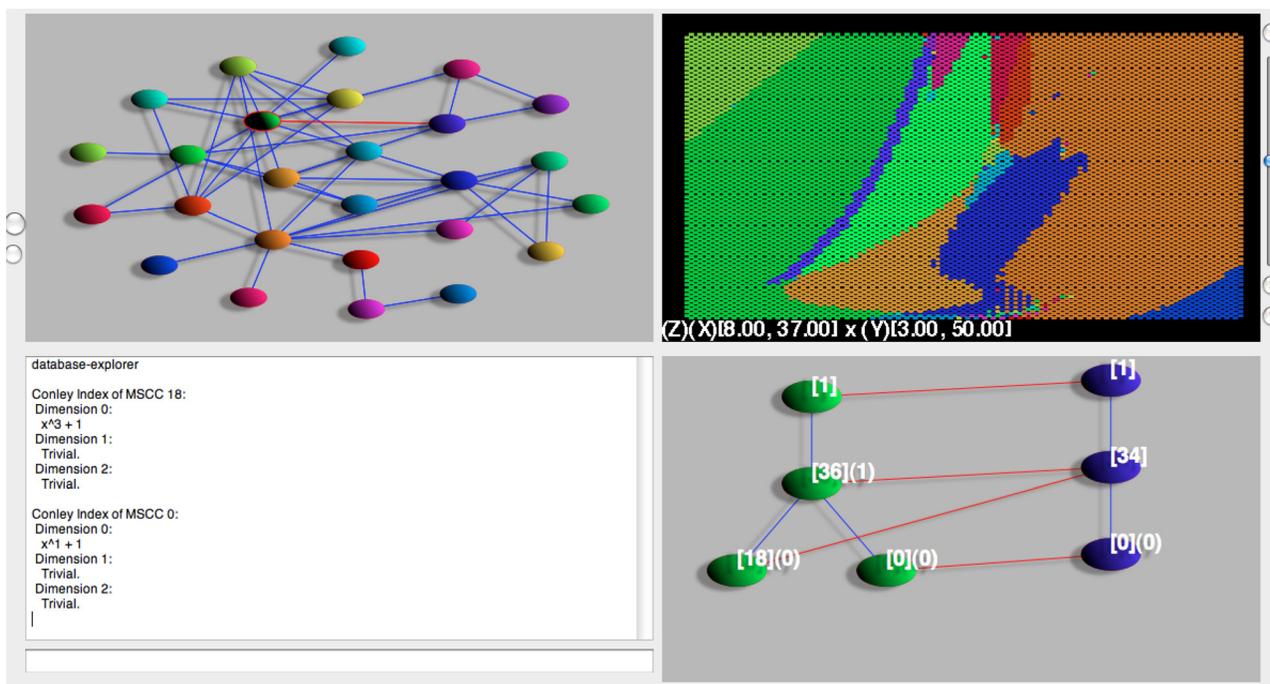


FIG. 2. Database information concerning dynamics for the overcompensatory Leslie model (5). (Upper left) Continuation Graph: Each node corresponds to a Conley-Morse graph equivalence class. Each edge corresponds to a clutching graph between Conley-Morse graphs. (Upper right) Parameter space divided into regions corresponding to Conley-Morse graph equivalence classes. Color coding of the nodes in the continuation graph. (Lower right) Clutching graph between two Conley-Morse graphs. This clutching graph corresponds to the highlighted (red) edge in the continuation graph. The nodes of the Conley-Morse graphs have two types of labels [*] and (*). The square brackets indicate the Morse continuation class associated with the node. The parenthesis indicate the level of homology on which the Conley index is nontrivial. (Lower Left) Conley indices of the Morse sets. The polynomial is the invariant factor for the shift equivalence class of the linear map on homology using \mathbb{Z}_2 coefficients induced by the index map. See Sec. III C.

refined parameter grid elements, the corresponding poset of the subgraph on a finer parameter grid element is projected onto the poset on a coarser parameter grid element by a map naturally induced by the parameter inclusion.

One of the problems that must be dealt with in the construction of Conley-Morse graphs is the appearance of Morse sets with trivial Conley index. Because of the trivial index, one cannot prove that their invariant part is nonempty, and in fact, for some systems, hundreds or thousands of such “spurious” sets isolating the empty set may appear. A powerful method for detecting such Morse sets in order to prune them and reduce the Morse graphs is to refine the grid in such a set and check for the emptiness of its invariant part at the higher resolution.² Therefore, if a Morse set does not appear to have the empty invariant part after several refinements then this is an indication of the presence of non-trivial dynamics and a possibility of bifurcation. See Ref. 1 for an analysis of the saddle-node bifurcation from this point of view.

The broader question of how continuation diagrams and Conley-Morse graphs relate to more classical approaches to dynamics like structural stability and bifurcations is a subtle one. In cases like the saddle-node bifurcation, the behavior of a Morse set under successive subdivisions can be suggestive of the underlying dynamics. In other cases, however, the coarse dynamics visible at a given level of subdivision provides very different information from what is known from the classical perspective.

As an example, consider the logistic map $f(x) = rx(1 - x)$ for $x \in [0,1]$. This map has an attracting fixed point

distinct from the origin for $r \in (1,3]$ and there is a well known period-doubling cascade for r between 3 and approximately 3.57. Figure 3 gives a schematic picture of the dynamics for $r \in [2.9, 3.7]$ on the level of the Morse graphs, where the phase space is subdivided at least 16 times (i.e., into 2^{16} boxes) and the interval $[2.9, 3.7]$ in parameter space is subdivided 10 times into 1024 boxes. (In this diagram, a small number of artifacts due to discretization are being ignored to convey the big picture.)

The Morse graphs in the diagram provide a schematic picture of the dynamics for the correspondingly labeled

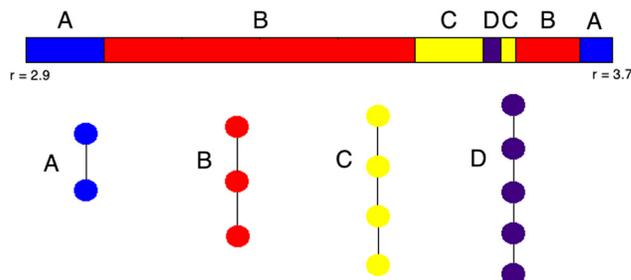


FIG. 3. Schematic picture of parameter space for the logistic map, $r \in [2.9,3.7]$. Here the transition from A to B occurs at 3.0078, B to C at 3.4563, and C to D at 3.5500. These transitions correspond to the first three period doublings, and they necessarily occur after the actual period doublings due to discretization in phase and parameter space. The transition from D back to C occurs at 3.5727, C to B at 3.5914, and B to A at 3.6781. These transitions do not correspond to bifurcations, but instead to the way the periodic attractor increasingly spreads out through the discretized phase space, making it impossible to distinguish from the unstable dynamics (except for the origin).

region of parameter space. This picture can be further refined by looking at the Conley index as in Sec. IV, though for now we will omit any details. It suffices to say that the Conley index verifies the existence of periodic behavior, as well as whether that behavior is stable or unstable.

The transition from region A to region B indicates the first period doubling bifurcation in the cascade. Likewise the transition from region B to region C corresponds to the second period doubling, and the transition from region C to region D the third. But, beyond this, further period doublings are invisible at this level of discretization. In fact, rather than the dynamics becoming finer, Figure 3 shows the dynamics becoming increasingly coarse as one proceeds through the period doubling cascade. These transitions do not correspond to bifurcations, but instead to the more global phenomenon of the periodic attractor moving around in a larger area of phase space given the level of discretization. Thus, as the period doublings accelerate, we actually see a sort of “mirror image” effect of the first few doublings. The important observation is that changes of continuation class can in some circumstances indicate the presence of classical bifurcations, but they can also be a consequence of dynamics without necessarily corresponding to any change in structural stability, and that the two possibilities cannot always immediately be distinguished.

III. ALGORITHMS

The mathematical framework proposed in Sec. II is combinatorial in nature and the presented motivation was the need for a robust description of dynamics with respect to both parameters and measurement. In this section, we change the perspective and observe that a combinatorial theory raises hopes that the dynamical structures being extracted are computable. Of course, the practicality of these computations depends upon the availability, development, and implementation of efficient algorithms.

A. Grid element container and combinatorial map structure

As is indicated in Sec. II A, we use grids and multivalued maps to translate between the continuous nonlinear world we are interested in and the combinatorial world we may compute in. This suggests the need for a programmatic interface to construct and interact with grids. In our approach, the grid elements themselves are identified with integers. The geometric objects of interest may be problem dependent and hence take the form of type *Geo*. Basic requirements are that the grid must provide methods for us to translate back and forth, as well as provide us access to a complete list of grid elements. Furthermore, we need to be able to adaptively subdivide these grids. This leads to the following constructs.

1. Iterator: A device used to loop through all grid elements in the grid \mathcal{X} .
2. Geometry: A method for producing the *Geo* object associated with a grid element.
3. Cover: A method for producing a set of grid elements which are known to provide an outer cover of a *Geo* object.

4. Subdivide: A method to change the structure of the grid by subdividing a grid element into smaller pieces.
5. Complex: A method to produce a cell complex out of a collection of grid elements.

For the computations presented in this paper, we use a grid based on multiscale cubical structures and deal exclusively with *Geo* objects which are rectangular prisms. This can be generalized without breaking the above interface.

Our representation of the dynamics is based on multivalued maps which are outer approximations. This cover produces enough grid elements to form an outer cover of the geometric region provided.

To capture the dynamics of F_ζ acting on $\mathbb{R}^n \times |\zeta|$ for a given $\zeta \in \mathcal{Z}$ requires a problem dependent computer routine \mathbb{F}_ζ that takes as input a *Geo* object A (think rectangular prism) and outputs a *Geo* object B , with the property that $F_\zeta(A, |\zeta|) \subset B$. The construction of \mathbb{F}_ζ makes use of a mixture of analytic bounds and rigorous numerical methods based on interval arithmetic (see for example Refs. 27 and 29). Typically, the choice of accuracy of \mathbb{F}_ζ involves a trade off between the difficulty of obtaining rigorous analytic bounds, the cost of evaluating such bounds, and memory constraints. With the routine \mathbb{F}_ζ , the multivalued map can be obtained by simple composition,

$$\mathcal{F}_\zeta = \text{cover} \circ \mathbb{F}_\zeta \circ \text{geometry}.$$

That is, we begin with a grid element, query the grid to recover a *Geo* object, apply the user-defined map function to produce another *Geo* object, and then cover this output with grid elements.

The grid provides a covering of the space in terms of topologically simple objects. However, to compute homology requires the finer structure of a cell complex. For simple grids consisting of cubes and simplices, the action of complex is classical. The use of more sophisticated grids requires the user to construct the appropriate complex operations.

B. Graph theory algorithms

For each $\zeta \in \mathcal{Z}$, the Morse graph MG_ζ is defined in terms of the multivalued map \mathcal{F}_ζ and provides a decomposition of the global dynamics. There are two essential computations associated with the construction of MG_ζ from \mathcal{F}_ζ : identification of the strongly connected path components and the partial order between these components which we refer to as *reachability*. Both memory and run time are critical issues that need to be addressed.

We begin by remarking that the computations are done in an adaptive manner; working with a uniform cubical grid decomposition of $X \subset \mathbb{R}^n$ is prohibitively expensive in both time and memory. A description of an adaptive procedure is presented in Ref. 2.

Tarjan's Algorithm is a standard procedure for computing strongly connected components.⁴ It proceeds by executing a depth first search and keeping track of so-called low-link information. Tarjan's algorithm requires time linear in the number of edges of the graph. It is important to note that storing \mathcal{F}_ζ can be memory intensive. The number of vertices

is related to the size and dimension of $\text{Inv}(X, F_\zeta)$ and the number of edges associated with each vertex is determined by the product of the eigenvalues of Df with magnitude greater than one. An approach to circumventing this problem is to avoid storing \mathcal{F}_ζ and instead using \mathbb{F}_ζ , recompute $\mathcal{F}_\zeta(\xi)$ whenever necessary. Naively running Tarjan’s algorithm with this approach leads to recomputing $\mathcal{F}_\zeta(\xi)$ many times for each grid element. Since these geometric computations can be quite expensive (especially in the context of differential equations), this is unacceptable. We remark that we have implemented a modified version of Tarjan’s algorithm such that we do not store the entirety of \mathcal{F}_ζ in memory and yet only have to evaluate $\mathcal{F}_\zeta(\xi)$ once for each $\xi \in \mathcal{X}$.¹⁰ More precisely, there exists a strongly connected components algorithm with space requirements that are linear in the number of vertices, but the run time is linear in the number of edges, and needs to query each vertex for its set of out-edges exactly once.

Turning to the issue of computing reachability, for a general directed graph there are no known linear time algorithms. However, we are not interested in the entire reachability relation, but only the reachability relation between Morse sets. While theoretically the number of Morse sets in a Morse decomposition can be arbitrarily large, in our experience this number is usually quite small. Because current computers deal in 64-bit words, it is possible to establish the reachability relation for up to 64 Morse sets in a single pass of the edges (which we process in topological order), i.e., in linear time. If there are more than 64 Morse sets, then we require multiple sweeps. We note that we have yet to see more than 64 Morse sets in an application.

C. Conley index algorithms

There are two issues associated with computing the Conley index. First, we need to be able to compute the map on homology induced by an index map. Second, we need to be able to identify shift equivalence classes on these maps. Addressing these questions in detail is beyond the scope of this paper, thus we provide a heuristic and cursory description of the challenges, how they are addressed and point the reader to appropriate references.

Recall from Proposition 2.8 that given a Morse set $\mathcal{M}_\zeta(p)$ associated with \mathcal{F}_ζ , we use the sets

$$\mathcal{P}_1 := \mathcal{F}_\zeta(\mathcal{M}_\zeta(p)) \quad \text{and} \quad \mathcal{P}_0 := \mathcal{F}_\zeta(\mathcal{M}_\zeta(p)) \setminus \mathcal{M}_\zeta(p)$$

to define an index pair $P = (P_1, P_0)$. The task is to compute

$$f_{P,*} : H_*(P_1/P_0, [P_0]) \rightarrow H_*(P_1/P_0, [P_0])$$

given $\mathcal{F}_\zeta, \mathcal{P}_1$, and \mathcal{P}_0 .

Our approach is to consider the graph of \mathcal{F}_ζ as a subset of $\mathcal{X} \times \mathcal{X}$. More precisely, we consider the pairs,

$$\begin{aligned} \mathcal{G}_1 &:= \{(\xi, \xi') \in \mathcal{P}_1 \times \mathcal{X} \mid \xi' \in \mathcal{F}_\zeta(\xi)\} \quad \text{and} \\ \mathcal{G}_0 &:= \{(\xi, \xi') \in \mathcal{P}_0 \times \mathcal{X} \mid \xi' \in \mathcal{F}_\zeta(\xi)\}. \end{aligned}$$

Working with a cubical grid it is clear that \mathcal{G}_1 and \mathcal{G}_0 generate cubical cell complexes¹⁴ on which we can define

projection maps π_d and π_r from the graph to the domain and range of \mathcal{F}_ζ , respectively. Under the assumption (which is satisfied for the examples presented in Sec. IV) that \mathcal{X} is the grid of a rectangular region in \mathbb{R}^n and \mathcal{F}_ζ is an outer approximation of a continuous function taking rectangular images it can be shown that

$$\pi_{d,*} : H_*(|\mathcal{G}_1|, |\mathcal{G}_0|) \rightarrow H_*(P_1, P_0) \cong H_*(P_1/P_0, [P_0])$$

is an isomorphism. Thus, in this setting it can be shown that

$$f_{P,*} := \pi_{r*} \circ \pi_{d*}^{-1} \tag{6}$$

provides an appropriate map on homology. The details of this construction in the more general setting of cell complexes can be found in Ref. 11.

The careful reader will note that $f_P : (P_1/P_0, [P_0]) \rightarrow (P_1/P_0, [P_0])$ is not well defined. What should be written is $f_{z,P}$, which is induced by f_z for some choice of $z \in \zeta$. However, the right hand side of Eq. (6) is well defined which implies that $f_{z,P,*}$ is independent of $z \in \zeta$. Note that this provides a local proof of N3 (compare with Refs. 3, 26, 23, and 8).

One challenge arises from the fact that the cell complex representing \mathcal{G}_1 can be quite large. Because of this, we have developed an algorithm that computes generators of $H_*(P_1, P_0)$ and lifts them to corresponding generators in $H_*(|\mathcal{G}_1|, |\mathcal{G}_0|)$, thus directly computing π_{d*}^{-1} . Furthermore, this procedure only needs to construct a single graph fiber at a time, and thus can have significantly lower memory requirements. Again the reader is referred to Ref. 11 for further details.

The computation of the shift equivalence class of $f_{P,k} : H_k(P_1/P_0, [P_0]) \rightarrow H_k(P_1/P_0, [P_0])$, which is the second aspect of computing the Conley Index, appears to be an open problem except in special cases. Fortunately, one such special case is that of vector spaces over finite fields. In this case, $f_{P,k}$ can be written as a matrix A , and the shift equivalence class of A can be determined from the rational canonical form of A . For a longer discussion of the rational canonical form, see Ref. 7.

Recall that the rational canonical form of square matrix A with entries in a field is a block diagonal matrix. The characteristic polynomial of each block is called an *invariant factor*. Two matrices are similar (in the usual linear algebra sense, meaning they represent the same linear transformation after a change of basis) if and only if the set of invariant factors for each matrix are identical, including multiplicity. In this way, a list of invariant factors is able to represent an equivalence class of similar matrices.

In the case of shift equivalence over a field, a modification of the invariant factors of A suffices to determine the shift equivalence class of A . The modification is easily described. To be concrete, assume the invariant factors are written as polynomials in x , where x is merely a formal parameter. Then, as a first step, divide through each invariant factor by the largest possible power of x , so that the resulting constant term is non-zero. Second, disregard all constant polynomials, i.e., those in which x does not appear. The resulting set of polynomials characterizes the shift equivalence class, in that A is shift

equivalent to another matrix B if and only if A and B are associated to the same set of polynomials, including multiplicity.

D. Continuation algorithms

As is indicated in Sec. II D, the concepts of continuation are based on the existence of clutching graphs. Computationally, we are presented with two Morse decompositions $\{\mathcal{M}_\zeta(p) \subset \mathcal{X} | p \in P_\zeta\}$ and $\{\mathcal{M}_\zeta(p') \subset \mathcal{X} | p' \in P_{\zeta'}\}$. Let n be the number of grid elements in all the Morse sets of both Morse decompositions. The following naive algorithm computes the clutching graph in $O(n^2)$ time. Use an outer loop which loops through every grid element in a Morse set in the first grid, and an inner loop that loops through the grid elements of the Morse sets of the second grid. Whenever an intersection is found, an edge in the clutching graph is forged. This naive algorithm, though generally applicable, is woefully inefficient. In practice, we re-express the Morse sets in one grid by covering them in the other. After this step, what remains is to scan grid elements, which takes linear time. Thus the complexity bottleneck is determined by how hard it is to cover a set of grid elements from $\cup_{p \in P_\zeta} \mathcal{M}_\zeta(p)$ with grid elements from $\cup_{p' \in P_{\zeta'}} \mathcal{M}_{\zeta'}(p')$, and vice versa. This, in turn, depends on the details of the grid implementation. For hierarchical tree-based multiscale cubical structures where the outer bounds of the grid are the same, we obtain an $O(n)$ algorithm.

Having determined the clutching graphs, it is easy to identify the Morse continuation classes and Morse Graph continuation classes via generating relations. However, we require a data structure which takes these generating equivalences as input and provide us with a representation of the equivalence classes. This is a classical problem and hence we employ the *disjoint set data structure*, also known as a *union-find* structure. This structure, when initialized, regards some finite set of elements as each belonging to disjoint singleton sets. By calling a union method, these sets may be unioned together until the disjoint set data structure represents the partition associated with the equivalence relation. (The find method is used to determining a representative element of each disjoint set; so it can be used to determine if two elements are equivalent). The union and find methods are not constant time, but rather the time complexity is given by the inverse Ackermann function.²⁸ For all practical purposes, we may consider inverse Ackermann to be constant time, as it grows *extremely* slowly. Given the union-find structure, what remains to us is simply to produce a generating set of relations to learn the equivalence classes. For each generating equivalence, we call the union method.

E. Database structure

The desired result of our computations is a database from which one can extract useful information concerning the dynamics. This information takes the form of a collection of *records*:

1. **Parameter Record:** Indicate a region of parameter space and give it an index so it may be cross-referenced by other records.
2. **Morse Record:** Indicate the Morse graph associated to some parameter record.
3. **Continuation Record:** Indicate the clutching graph associated to some indicated pair of parameter records.
4. **Conley Record:** Indicate the Conley index associated with some indicated Morse set associated to some indicated Morse Record.

If we only have records of type 1–3, we call it a *Morse Graph database*. If, on the other hand, we have all types of records, it is a *Conley-Morse Graph database*. In our computations, we first produce a Morse Graph database, and then process the continuation records using a union-find structure in order to learn the Morse continuation classes. Then we choose a single representative of that class, and compute the “Conley Record” associated with it. By Corollary 2.16 the Conley index is constant on continuation classes of Morse sets. Clearly, this is much more efficient than performing expensive Conley Index computations to produce Conley Records for every Morse set of every Morse Record.

F. Query algorithms

Once we obtain the database structure, we make use of it via *database queries*. Because the records make reference of each other, it is possible to make a number of different queries. Examples include, but are not limited to:

1. Identify all parameter regions in the same Morse Graph continuation class.
2. Identify all Morse Graph continuation classes for which we have shown multiple basins of attraction exist.
3. Identify all Morse Graph continuation classes which contain a Morse set with a given Conley index.

IV. APPLICATIONS

In order to demonstrate the potential of our approach, we present two applications of the computational framework introduced in the previous sections.

A. Three age-classes overcompensatory Leslie model

Consider the following three-age class overcompensatory Leslie population model,

$$\begin{cases} x = (\theta_1 x + \theta_2 y + \theta_3 z) e^{-0.1(x+y+z)} \\ y = 0.7x \\ z = 0.7y \end{cases} . \quad (7)$$

The variables x , y , and z represent the age class populations in order of ascending age. If one views this as a model of a plant population, then the 3 parameters θ_i , $i = 1, 2, 3$ can be interpreted as the seed production rates of the different age classes. The exponential term represents an overcrowding factor that depends on the adult population. This model and

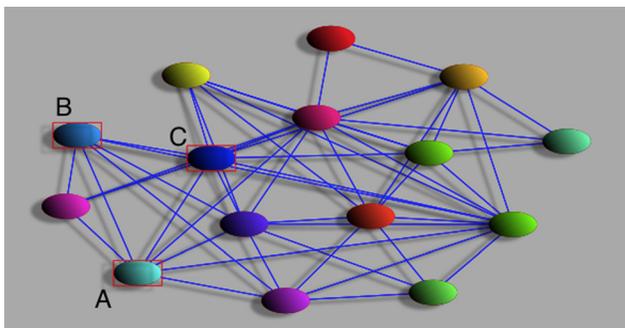


FIG. 4. Continuation graph showing the 15 largest continuation classes for the system (7). The boxed nodes represent Morse graphs with multiple minimal nodes which implies that at the corresponding parameter values there exist multiple basins of attraction. The class labeled A contains 13 964 parameter boxes, B contains 5222 boxes, and C contains 1497 boxes.

its biological relevance are discussed in greater detail in Ref. 30.

To set up the computations, we choose the parameter region of interest,

$$Z := \{(\theta_1, \theta_2, \theta_3) | 14.5 \leq \theta_1 \leq 30.5, 13.0 \leq \theta_2 \leq 37.0, 13.0 \leq \theta_3 \leq 37.0\} \subset \mathbb{R}^3.$$

The parameter space grid Z is constructed by subdividing Z into 32 equal sized intervals in each direction. This divides parameter space into 32 768 three dimensional cubical cells. To speed up the computation, we choose to compute for sets of parameter values that are represented by the edges in this complex. Thus, there are a total of 104 544 one dimensional parameter boxes for which we produce a Conley-Morse graph.

The compact region in phase space is given by

$$X := [0, 320.056] \times [0, 224.040] \times [0, 224.040] \subset [0, \infty)^3.$$

It can be shown analytically that X is an isolating neighborhood for the global attractor of Eq. (7) for all $\theta \in Z$. Using an adaptive subdivision algorithm similar to that of Ref. 2, the final grid in phase space consists of cubes obtained by subdividing X into 2^{12} equal sized intervals in each direction.

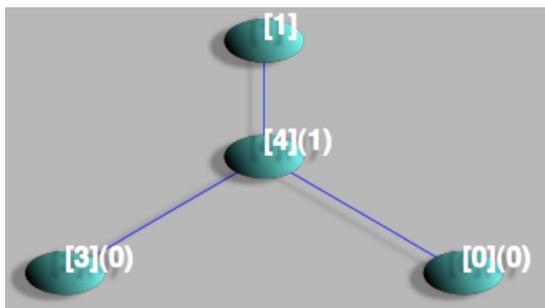


FIG. 5. Conley-Morse graph for continuation class A in Figure 4. Here the partial ordering between two nodes joined by an edge is indicated by their relative heights (the lower node being less than the higher node). The bracketed numbers on each node indicate the Morse set continuation class number, and the numbers in parentheses indicate any dimensions in which the Conley index is nontrivial. Here MSCC[1] contains the unstable origin. MSCC[4] is also unstable, as indicated by the nontrivial Conley index on level 1. The Morse graph further suggests saddle-like behavior separating the two attractors, although more analysis is required to prove this.

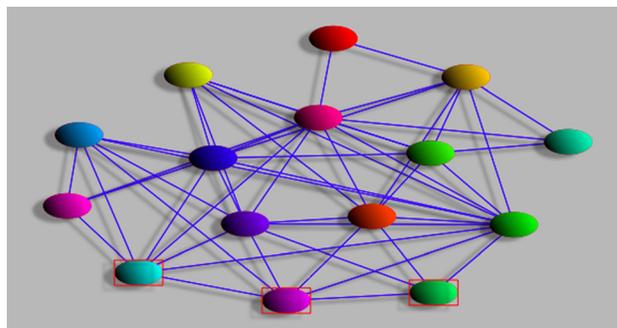


FIG. 6. Three boxed nodes representing the continuation classes over which the Morse set continuation class of MSCC[0] in Figure 5 extends. By the discussion in the text, MSCC[0] is an attracting 1-cycle set. Although the three nodes represent three continuation classes, hence potentially different dynamics, any difference in the observed dynamics must occur away from MSCC[0].

In general, this produces Morse decompositions consisting of many Morse sets most of which are spurious in the sense that the recurrence is due to coarseness in the discrete representation of phase space. To eliminate obvious spurious solutions, the cubes in each Morse set are once again subdivided up to 2^4 times in each direction and recurrence within these regions is checked. If the recurrence disappears, then one can conclude that the associated region in phase space does not contain recurrent dynamics. For a more detailed discussion of this step, see Ref. 2.

The computation based on the aforementioned inputs was run on 15 nodes of a cluster, using 3 processors per node. Each node had a minimum of 8 GB of memory. The total computation time was 137h, of which 134h were needed to find the Morse graphs.

As indicated in "Introduction," an important impetus behind the database is to be able to quickly and efficiently find interesting dynamics. Given that this is a population model, a natural question concerns the structure of attractors and/or the existence of multiple attractors as a function of the parameter values.

Let us begin by considering the existence of multiple attractors. The appropriate query is to ask for those nodes in the continuation graph for which the associated Morse graphs contain more than one minimal node. To make the results visually manageable, we restrict our attention to the larger continuation classes. In particular, Figure 4 shows the

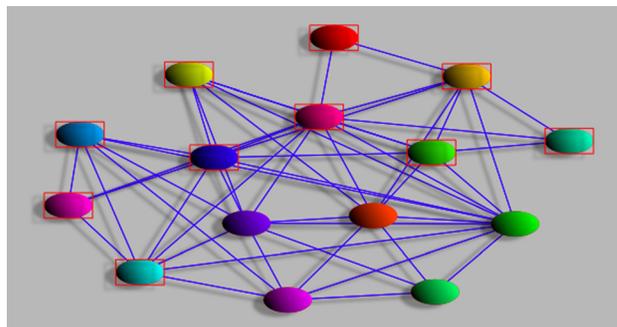


FIG. 7. Ten boxed nodes representing the continuation classes over which the Morse set continuation class of MSCC[3] in Figure 5 extends. One can conclude that MSCC[3] is a 4-cycle set from the Conley index information. Therefore all of the boxed nodes indicate regions of parameter space exhibiting this behavior.

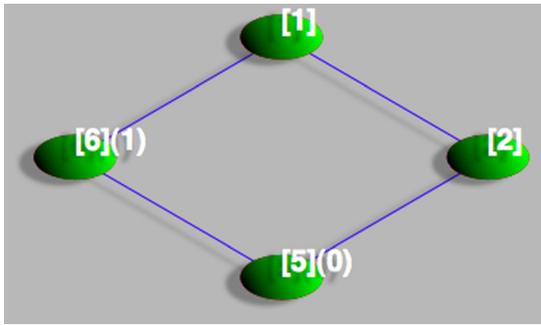


FIG. 8. A Conley-Morse graph with an attractor that belongs to the Morse set continuation class MSCC[5]. Here MSCC[5] is a 2-cycle set, and the only attractor since it is the only minimal node. MSCC[6] has nontrivial Conley index on the first level, which proves it contains recurrent dynamics. Since the only Morse set below MSCC[6] in the partial order is MSCC[5] there must be a connecting orbit between the two.

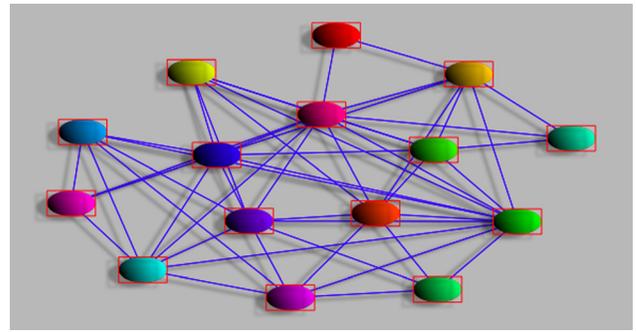


FIG. 9. Boxed nodes indicate continuation classes for which the Morse set containing the origin is not an attractor. Biologically, these can be considered regions that exhibit persistence or permanence, i.e., regions where the attractor is separated from the origin (extinction) at the scale of the computation. Here, all displayed nodes are boxed, so one can conclude that most of parameter space exhibits persistence.

continuation graph for the 15 largest continuation classes. These 15 continuation classes are associated with 103 593 of the 104 544 grid elements of parameter space which implies that over 99% of parameter space is accounted for. Recall that two nodes are adjacent in this graph if they represent adjacent regions of parameter space, so that this gives, in a sense, a schematic picture of the three-dimensional parameter space.

The boxed nodes in Figure 4 represent Morse graphs with multiple minimal nodes and therefore for the corre-

sponding parameter values, there are multiple basins of attraction. Observe that there are three mutually adjacent continuation classes with multiple attractors. The sizes of these parameter regions are 13 964, 5222, and 1497 parameter boxes. Thus in roughly 20% of the parameter space we study, we can guarantee the existence of at least two basins of attraction.

We now turn to the question of identifying the structure of the dynamics in the attractors. In particular, we make use of the following language.

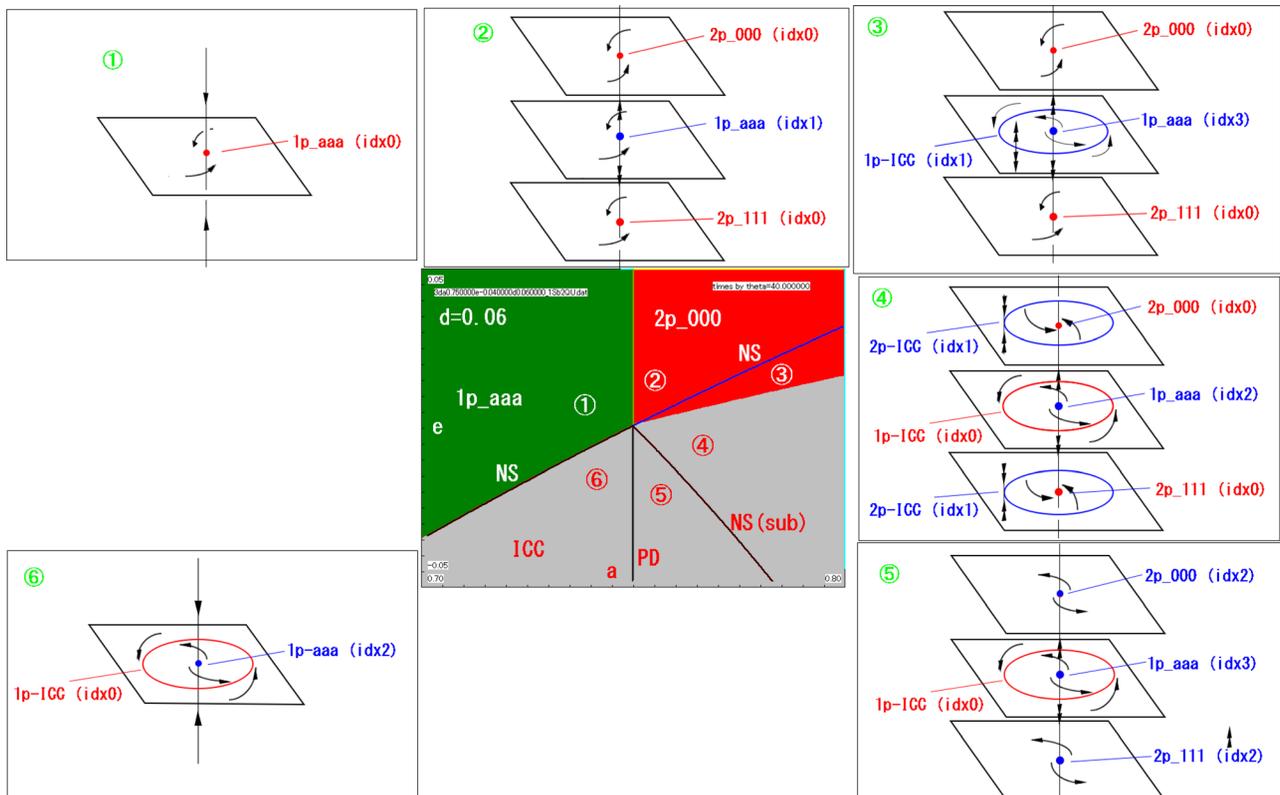


FIG. 10. The numerically generated bifurcation diagram (colored plate in the center) and corresponding schematic phase portraits for 3-CML, obtained by Komuro. In the bifurcation diagram in the center, the following abbreviations are used: PD = period-doubling bifurcation, NS = (supercritical) Neimark-Sacker bifurcation, NS(sub) = subcritical Neimark-Sacker bifurcation, ICC = invariant closed circle. The parameter range is chosen to be $[0.7, 0.8] \times [-0.05, 0.05]$ with $\delta = 0.06$, which is divided into six regions, numbered by circled numerals. Each parameter region corresponds to the plate numbered by the same circled numeral. The invariant closed circles appear in Regions 3–6. Especially, in Region 4, there exist multiple invariant closed circles, the unstable one being periodic with period two under the action of 3-CML. Figure courtesy of M. Komuro.

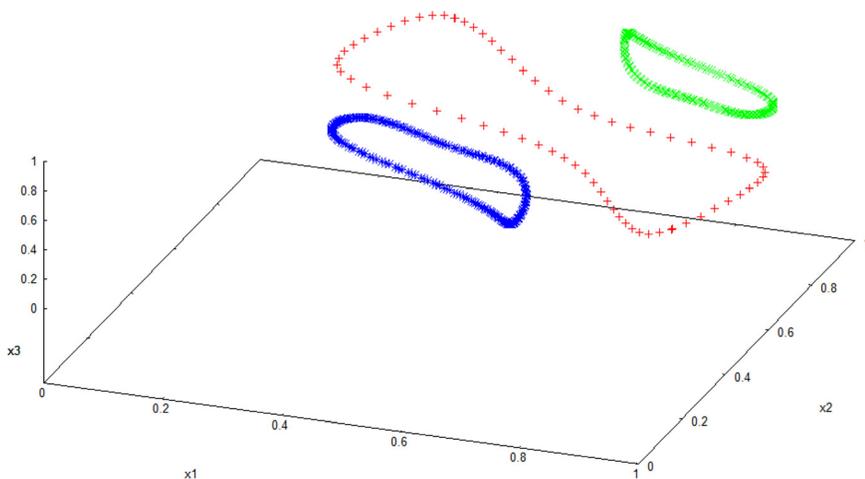


FIG. 11. The invariant circles appearing in Region 4 in Figure 10, computed numerically by M. Komuro. The red circle is stable, while the green and blue ones are unstable. Figure courtesy of M. Komuro.

Definition 4.1. An isolated invariant set S for a map $f : X \rightarrow X$ is a T -cycle set if there exist T disjoint, compact regions N_1, \dots, N_T such that $S = \text{Inv}(N, f)$ where $N := \bigcup_{i=1}^T N_i$ is an isolating neighborhood, and

$$f(N_i \cap N) \subset N_{i+1}, i = 0, \dots, T - 1,$$

where $N_0 := N_T$.

Consider the continuation class A with 13 964 parameter boxes. The associated Conley-Morse graph is shown in Figure 5. Recall that the bracketed numbers identify the Morse set continuation class (MSCC) of each node in the graph. The Conley index of MSCC[0] is trivial except in dimension zero, where it has invariant factor $x + 1 \pmod{2}$. By [Ref. 2, Proposition 5.8], we conclude that the associated Morse set is a 1-cycle set. More specifically there is an associated isolating neighborhood which is contractible and maps strictly to its interior under the dynamics of Eq. (7).

As is indicated in Sec. IID, the above mentioned description in terms of the 1-cycle set extend to the entire Morse set continuation class MSCC[0]. The three boxed nodes in Figure 6 indicate Morse set continuation classes over which MSCC[0] extends.

The other attractor, MSCC[3], has Conley index which is trivial except in dimension zero where it is represented by the invariant factor $x^4 + 1 \pmod{2}$. This is characteristic of a 4-cycle set. The extent of this Morse set continuation class is given by the boxed nodes in the continuation graph in Figure 7.

Observe that the nodes in the continuation graph, which are not boxed in either Figure 6 or Figure 7, must contain another distinct attracting Morse set continuation class. It appears, for example, as MSCC[5] in the Conley-Morse graph shown in Figure 8. The Conley index of MSCC[5] in dimension zero is represented by $x^2 + 1$, which is characteristic of a 2-cycle set. Among the fifteen largest continuation classes, these are all of the Morse set continuation classes of attractors that appear.

Observe that we have characterized the attractors for a large fraction of parameter space. However, because of the concern of extinction in the context of small perturbations biologists often are interested in understanding when the attractor is bounded away from the states of extinction. This is often called persistence or permanence (see Ref. 13 for a

precise definition and further discussion). In the setting of this model, extinction can be identified with the origin. We remark that the origin appears in the database as MSCC[1]. It can also be checked that MSCC[1] extends over all of the fifteen largest continuation classes (see Figure 9). Furthermore, in none of these cases, it is an attracting Morse set (special cases of this can be seen in Figures 5 and 8). Since the computations have been done at a minimal fixed scale, we can conclude that Eq. (7) exhibits persistence.

Of course there are additional questions that can be asked concerning the dynamics of Eq. (7). For some of these, additional database queries can prove useful. However, we hope that in the context of attractors, which is the most reasonable entry point for questions concerning biologically observable phenomena, we have made it clear that the database provides sufficient information to obtain useful nontrivial answers.

B. Three-dimensional coupled map lattice

As another example, the Conley-Morse Database software has been applied to the coupled map lattice (CML), a coupled system of maps introduced by Kaneko and others independently around early 1980s, see Ref. 16 and references therein for more information.

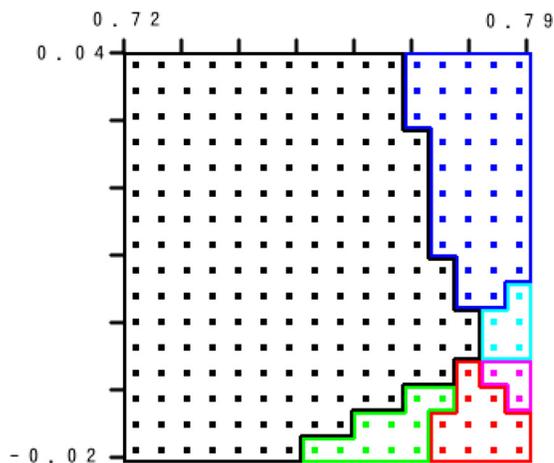


FIG. 12. The continuation diagram computed by the Conley-Morse database software. Compared to Figure 10, the number of divided regions of the parameter domain agrees, but each region is a little tweaked, which is mainly due to the overestimation by the construction of multivalued maps.

The n -CML is an n -dimensional dynamical system $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$, with $F(x) = (F_1(x), \dots, F_n(x))$ for $x = (x_1, \dots, x_n)$ given by

$$F_i(x) = (1 - \epsilon)f_a(x_i) + \frac{1}{2}(\epsilon - \delta)f_a(x_{i-1}) + \frac{1}{2}(\epsilon + \delta)f_a(x_{i+1}),$$

$$i = 1, \dots, n,$$

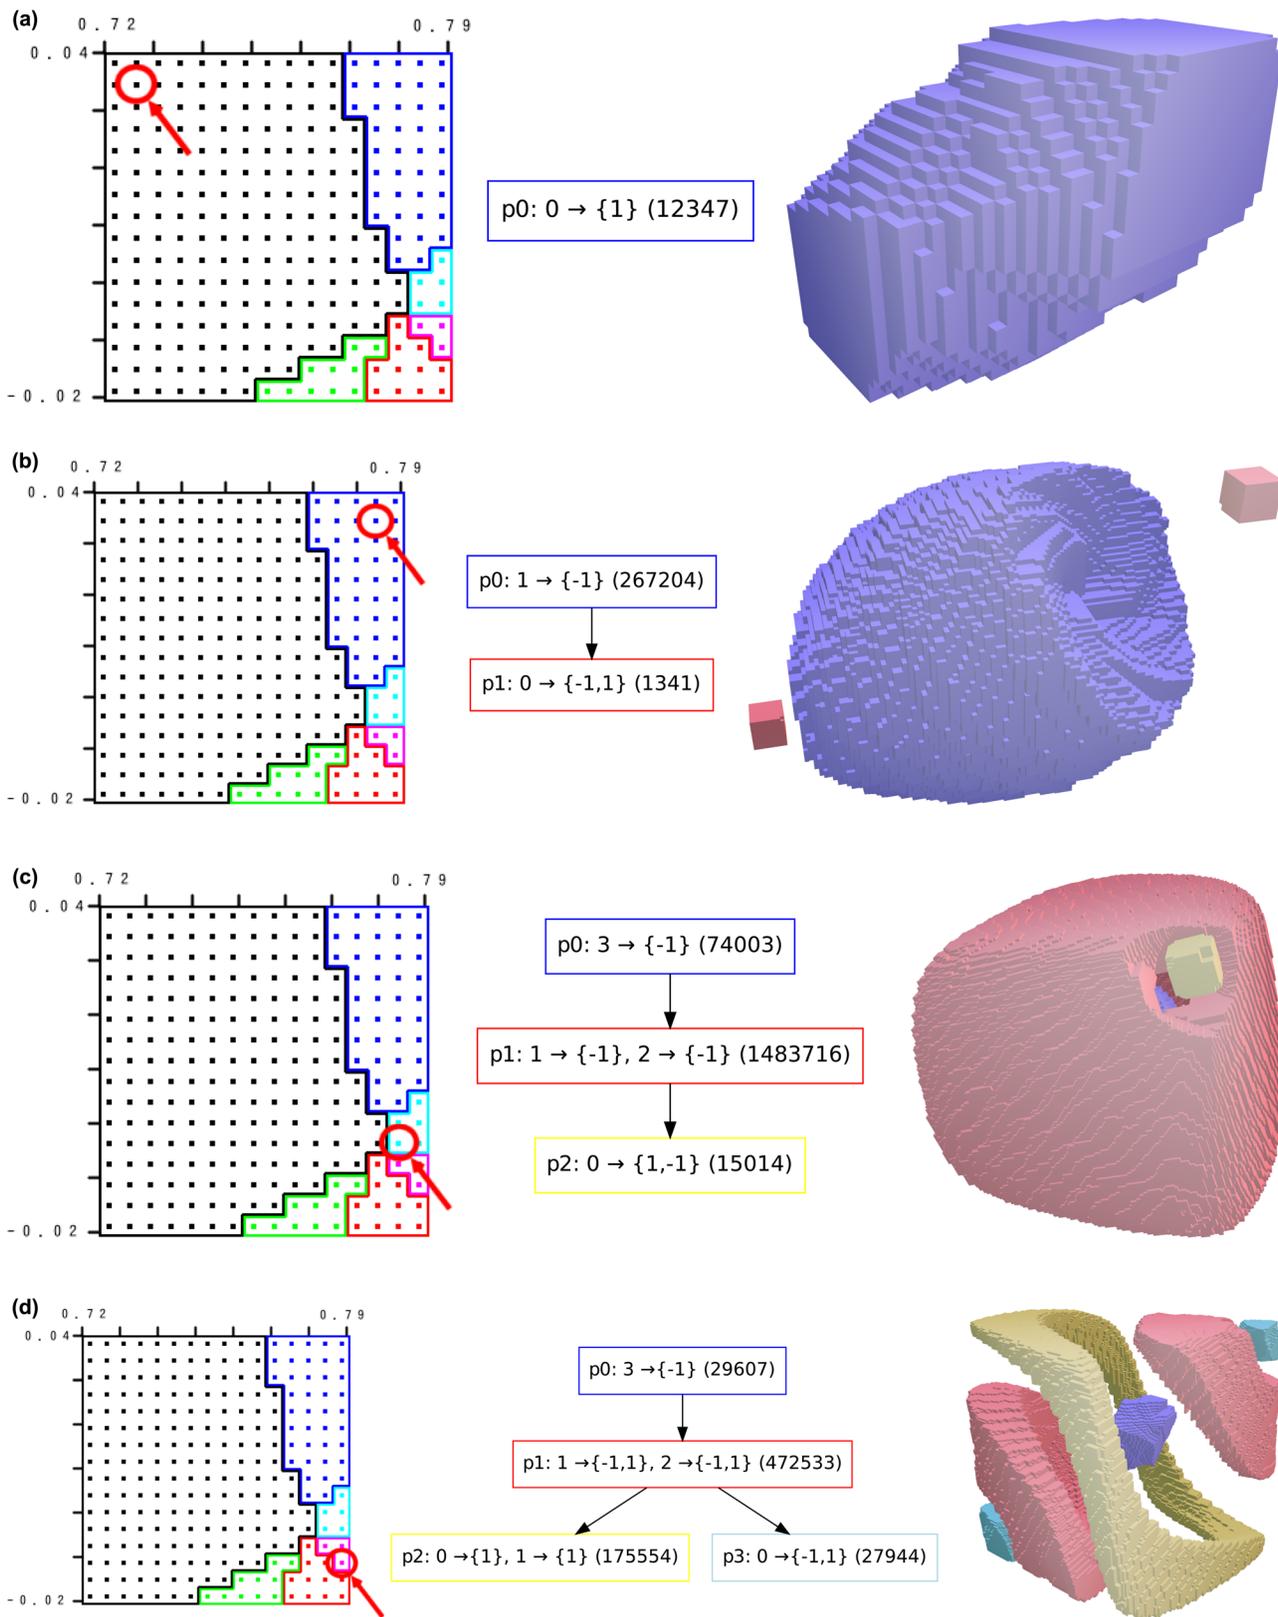


FIG. 13. Results of the computations for the 3-CML. For each Morse graph continuation class, a Conley-Morse graph and the corresponding Morse sets are plotted for a sample parameter point. The choice of the parameter for each class is marked with a red circle at the left-hand side, the Conley-Morse graph for the parameter is drawn in the center, and the Morse sets for the parameter appear at the right-hand side.

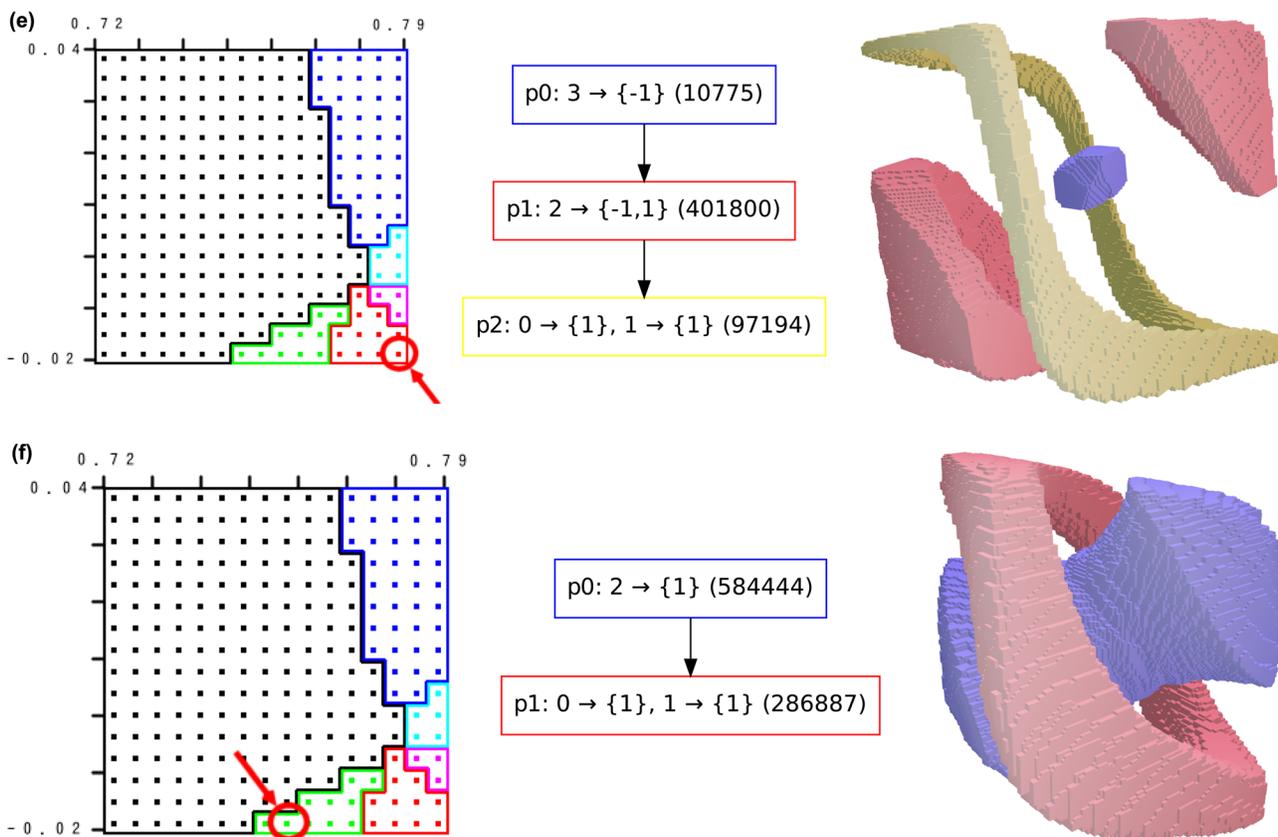


FIG. 13. (Continued).

where $x_0 = x_n$ and $x_{n+1} = x_1$. This system has three parameters, a , ε , and δ . In this paper, we choose f_a to be the logistic map, $f_a(\xi) = 1 - a\xi^2$.

There is a well-studied, coupled dynamical system, similar to CML, known as the globally coupled maps (GCMs)¹⁷ defined by

$$F(x)_i = (1 - \varepsilon)f_a(x_i) + \frac{\varepsilon}{n} \sum_{j=1}^n f_a(x_j) \quad i = 1, \dots, n.$$

As an important feature of CML, as compared to GCM, CML is a non-symmetric coupling, and as a result, CML exhibits a traveling wave at some parameters. In order to study the traveling waves more closely, Komuro³¹ examined the 3-CML in detail, using conventional numerical analysis. The object corresponding to traveling waves in 3-CML is an invariant closed circle (ICC), which can be observed numerically in a region of the (a, ε) -parameter plane of 3-CML with fixed $\delta = 0.06$. See Figures 10 and 11 for a numerically generated bifurcation diagram obtained by Komuro using the numerical method proposed in Ref. 18.

Below, we describe the result of computation of the 3-CML using the Conley-Morse Database software. In fact, the actual computation was done using the first version of the software explained in Ref. 2, as the latest version of the software discussed in the previous sections was not yet fully available at the time of the computation.

The computation parameters are taken as follows:

- (a, ε) varies in $[0.72, 0.79] \times [-0.02, 0.04]$
- δ is fixed to 0.06.

Note that in this parameter region, f_a has attracting period-2 periodic points. The box of $[0.72, 0.79] \times [-0.02, 0.04]$ is divided into 16×16 small boxes, and we take the center of each box as the input parameter value for each computation. As the interval arithmetic is used in the software, we could have taken the entire small boxes as the input parameters for the computations, but this would have been much more time consuming, and the overestimates would require us to use a finer grid in the parameter space. The phase space for the logistic map is taken as $[-1.1, 1.1]$, hence the entire phase space is $[-1.1, 1.1]^3$, on which we put the uniform grid decomposition into $2^9 \times 2^9 \times 2^9$ boxes. When needed, we set the refinement level of subdivision as 4, in case the computation result is not sufficient.

The Conley-Morse graph and corresponding phase space structure for each parameter region in Figure 12 is shown in Figure 13. The format of the Conley index information of the Conley Morse graphs follows (Ref. 2). For instance, the expression $p0:0 \rightarrow \{1\}(12347)$ in the Conley-Morse graph in Region No. 1 indicates that there is a (unique in this case) Morse set labelled p0 consists of 12 347 cubes in this computation, and the corresponding Conley index is non-trivial in the 0-th level of homology, whose associated homology map has non-zero eigenvalue 1. Similarly, the Conley-Morse graph in Region No. 3 contains three Morse sets, and the largest one labelled p1, identified by the frame color to see that it corresponds to the pink set in the right figure, consists of 1 483 716 cubes, and the corresponding Conley index is non-trivial in the 1st and 2nd levels of homology, with both of the homology maps having non-zero eigenvalue -1 . These

results agree well with the bifurcation structure shown in Figure 10.

ACKNOWLEDGMENTS

M.G. acknowledges the support of Fapesp Processo 2010/00875-9 and CNPq Processo 306453/2009-6. H.K. acknowledges the support of Ministry of Education, Science, Technology, Culture and Sports, Japan, Grant-in-Aid for Scientific Research No. 21340035. The work of K.M., S.H., and J.B. was partially supported by NSF Grants DMS-0915019 and CBI-0835621 and by contracts from DARPA and AFOSR. P.P. acknowledges the support from the E.U. and Portuguese national funds received through Fundação para a Ciência e a Tecnologia (FCT), project FCOMP-01-0124-FEDER-010645 (Ref. FCT PTDC/MAT/098871/2008) and Est-C/MAT/UI0013/2011.

¹Z. Arai, M. Gameiro, T. Gedeon, H. Kokubu, K. Mischaikow, and H. Oka, "Graph-based topological approximation of saddle-node bifurcation in maps," *RIMS Kokyuroku Bessatsu* **B31**, 225–241 (2012).

²Z. Arai, W. Kalies, H. Kokubu, K. Mischaikow, H. Oka, and P. Pilarczyk, "A database schema for the analysis of global dynamics of multiparameter systems," *SIAM J. Appl. Dyn. Syst.* **8**, 757–789 (2009).

³C. Conley, "Isolated invariant sets and the Morse index", in *CMBS Regional Conference Series in Mathematics* (American Mathematical Society, 1978), Vol. 38.

⁴T. Cormen, C. Leiserson, and R. Rivest, *Introduction to Algorithms*, The MIT Electrical Engineering and Computer Science Series (MIT, Cambridge, MA, 1990).

⁵S. Day, O. Junge, and K. Mischaikow, "A rigorous numerical method for the global analysis of infinite-dimensional discrete dynamical systems," *SIAM J. Appl. Dyn. Syst.* **3**, 117–160 (2004).

⁶S. Day, Y. Hiraoka, K. Mischaikow, and T. Ogawa, "Rigorous numerics for global dynamics: A study of the Swift-Hohenberg equation," *SIAM J. Appl. Dyn. Syst.* **4**, 1–31 (2005).

⁷D. S. Dummit and R.M. Foote, *Abstract Algebra*, 3rd ed. (Wiley, 2004).

⁸J. Franks and D. Richeson, "Shift equivalence and the Conley index," *Trans. Am. Math. Soc.* **352**, 3305–3322 (2000).

⁹M. Gameiro, T. Gedeon, W. D. Kalies, H. Kokubu, K. Mischaikow, and H. Oka, "Topological horseshoes of traveling waves for a fast-slow predator-prey system," *J. Dyn. Differ. Equ.* **19**, 623–654 (2007).

¹⁰S. Harker and K. Mischaikow, "Topological sort for graphs with large numbers of edges" (in preparation).

¹¹S. Harker, K. Mischaikow, M. Mrozek, and V. Nanda, "Discrete Morse theoretic algorithms for computing homology of complexes and maps," *Foundations of Computational Mathematics* (accepted).

¹²A. Hatcher, *Algebraic Topology* (Cambridge University Press, 2002).

¹³V. Hutson and K. Schmitt, "Permanence and the dynamics of biological systems," *Math. Biosci.* **111**, 1–71 (1992).

¹⁴T. Kaczynski, K. Mischaikow, and M. Mrozek, *Computational homology*, in *Applied Mathematical Sciences* (Springer-Verlag, New York, 2004), vol. 157.

¹⁵W. D. Kalies, K. Mischaikow, and R. C. A. M. Vandervorst, "An algorithmic approach to chain recurrence," *Found. Comput. Math.* **5**, 409–449 (2005).

¹⁶K. Kaneko, "Overview of coupled map lattices," *Chaos* **2**, 279–282 (1992).

¹⁷K. Kaneko, "Clustering, coding, switching, hierarchical ordering, and control in a network of chaotic elements," *Physica D* **41**, 137–172 (1990).

¹⁸K. Kamiyama, M. Komuro, and T. Endo, "Bifurcation of quasi-periodic oscillations in mutually coupled hard-type oscillators—Demonstration of unstable quasi-periodic orbits," *Int. J. Bifurcation Chaos* **22**, 1230022 (2012).

¹⁹K. Mischaikow, M. Mrozek, J. Reiss, and A. Szymczak, "Construction of symbolic dynamics from experimental time series," *Phys. Rev. Lett.* **82**, 1144–1147 (1999).

²⁰K. Mischaikow and M. Mrozek, "Conley index," in *Handbook of Dynamical Systems* (North-Holland, Amsterdam, 2002), vol. 2, pp. 393–460.

²¹K. Mischaikow and M. Mrozek, "Chaos in the Lorenz equations: A computer-assisted proof," *Bull. Amer. Math. Soc.* **32**, 66–72 (1995).

²²K. Mischaikow, M. Mrozek, and P. Pilarczyk, "Graph approach to the computation of the homology of continuous maps," *Found. Comput. Math.* **5**, 199–229 (2005).

²³M. Mrozek, "Leray Functor and the cohomological Conley index for discrete time dynamical systems," *Trans. Am. Math. Soc.* **318**, 149–178 (1990).

²⁴M. Mrozek, "An algorithm approach to the Conley index theory," *J. Dynam. Differ. Equ.* **11**, 711–734 (1999).

²⁵J. Munkres, *Topology*, 2nd ed. (Prentice Hall, 2000).

²⁶J. W. Robbin and D. Salamon, "Dynamical systems, shape theory and the Conley index," *Ergod. Theory Dyn. Syst.* **8**, 375–393 (1988).

²⁷S. Rump, "Verification methods: rigorous results using floating-point arithmetic," *Acta Numerica* **19**, 287–449 (2010).

²⁸R. E. Tarjan and J. van Leeuwen, "Worst-case analysis of set union algorithms," *J. ACM* **31**(2), 245–281 (1984).

²⁹W. Tucker, *Validated Numerics: A Short Introduction to Rigorous Computations* (Princeton University Press, 2011).

³⁰I. Ugarcovici and H. Weiss, "Chaotic dynamics of a nonlinear density dependent population model," *Nonlinearity* **17**, 1689–1711 (2004).

³¹M. Komuro, personal communication 2011.