

Chapter 1

How topology came to chaos

Robert Gilmore

aindx]Gilmore, Robert

*Physics Department
Drexel University
Philadelphia, PA 19104
USA*

We review the steps taken during the development of topological analysis tools for the analysis of chaotic data.

Contents

1. How topology came to chaos 1
Robert Gilmore

1. Background

Feigenbaum's startling and initially unappreciated discovery of the universality of certain scaling ratios¹⁻³ set off a chain of events that continues to this day. Before this discovery it had been assumed that

All linear systems are the same.
Each nonlinear system is nonlinear in its own way.^{4,5}

After this discovery it was realized that

It was a very happy and shocking discovery that there were structures in nonlinear systems that are always the same if you looked at them in the right way.^{4,6}

Once a theory has been developed to the point where predictions are possible, there is a rush to experiments to falsify or to confirm the theory. Such was the case with Feigenbaum's scaling predictions. Experiments were carried out on fluids,^{7,8} chemical reactions,⁹⁻¹¹ electrical circuits,¹² and lasers.¹³⁻¹⁶ The experiments differed in the physics involved and the time scales involved, ranging from months (fluid experiments) to days (chemical reactions) to minutes (electrical circuits) to

milleseconds (laser experiments). In the end, all confirmed Feigenbaum's predictions of the universality of scaling relations according to one-dimensional maps of the interval.

On the one hand the return on the investment of time required to do these experiments was magnificent: the experiments had shown that there were at least three previously unknown constants of Nature — scaling ratios in state space ($\alpha = -2.50290\dots$), in the control parameter space ($\delta = +4.466920\dots$), and in the space of measured intensity ratios (-8.2 dB). On the other hand the results were disappointing. In the logistic map of the interval, beyond the initial period-doubling bifurcation where the scaling relations are predicted, as well as all subsequent windows, there is a rigid organization of behavior — an organization of periodic windows in a sea of chaotic behavior. Indications of the rigidity of this structure were seen and reported in many of the experiments. But the nature and the structure of this rigidity was yet to be determined.

2. A Challenge

At this point Prof. J. R. Tredicce offered a challenge — and an opportunity. He had a great deal of data “left over” from his “Feigenbaum experiments” carried out on a laser with modulated losses.¹⁶ Since data are always acquired at high cost with pain and difficulty, it would be obscene not to make further use of them. He asked if I could help him to further “understand” these data. This seemed like a wonderful challenge at the time. And with the benefit of hindsight, it was even more than that.

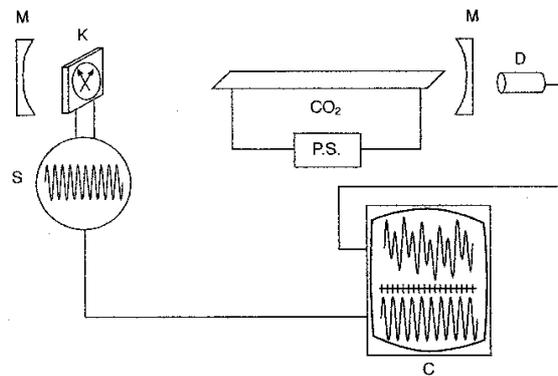


Fig. 1. Schematic diagram of a laser with modulated losses.¹⁶ A carbon dioxide gas tube (CO₂, P.S. is the power source) is inserted between two mirrors (M) that form a confocal resonant cavity. A Kerr cell (K) is inserted in the cavity. This cell is periodically modulated by a signal (S), inducing losses as the polarization of the cell deviates from that produced by the Brewster angle windows. The intensity output is recorded by the detector (D). The input and output signals are recorded in a computer (C).

At the time there were only two types of tools available for the study of experimental data from a chaotic source. One type depended on metric measures: fractal dimension estimates of all types.^{17–20} The other type depended on dynamical measures: Lyapunov exponents and spectra of Lyapunov exponents.^{21–25} Both required very long, very clean data sets, a great deal of computation, resulted in real number estimates with no realistic error estimates,²⁶ both were often dependent on where and how to make certain crucial assumptions, and neither was generally reproducible. Further, neither type of tool provided a way to distinguish among the different types of strange attractors that could be seen to be different: for example the mathematical strange attractors of Lorenz²⁷ and Rössler,²⁸ not to mention the strange attractors associated with periodically driven nonlinear oscillators such as the Duffing and van der Pol oscillators.

It was clear that a new type of analysis methodology was called for. It was also clear that this new tool should not depend on metric invariants or dynamical invariants. In fact, looking back to Poincaré,²⁹ it was clear that this new tool ought to be topological in nature. Listening more closely to Poincaré, it was clear that this new tool ought to involve the periodic orbits “in” a chaotic attractor. A chaotic trajectory winds around in phase space arbitrarily close to any unstable periodic orbit, so it ought to be possible to use segments of a chaotic trajectory as good approximations (surrogates) for unstable periodic orbits.^{30,31} The location and identification of such orbits is sometimes simplified because many unstable periodic orbits first appear in their stable avatars when created in saddle-node or period-doubling bifurcations and don’t move too far from their original position in phase space as control parameters are changed. It was clear that unstable periodic orbits could not only be extracted from chaotic time series³² but also serve as the “skeleton” of the strange attractor.³⁰ This is illustrated in Fig. 2.

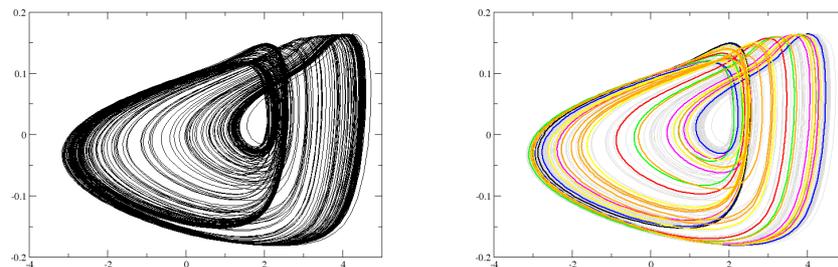


Fig. 2. (color online) The strange attractor (left) that describes the Belousov-Zhabotinskii reaction is very well outlined by its skeleton, a superposition of unstable periodic orbits of low period (right).

3. Relative Rotation Rates

It was our hope that periodic orbits would somehow provide a key to understanding the structure of chaotic attractors. With this idea in mind, Hernan G. Solari extracted a number of unstable periodic orbits from a mathematical model of Tredicce's laser with modulated losses and undertook to determine the topological properties of these orbits.³³ The simplest and most obvious tool for quantitatively understanding the topological organization of these periodic orbits was to consider them as oriented closed loops and compute their Gauss linking numbers. After all, the origin of Gauss linking numbers came from two different branches of physics. The linking number of two oriented closed loops, or periodic orbits A and B , is defined by the Gauss integral

$$LN(A, B) = \frac{1}{4\pi} \oint \oint \frac{(\mathbf{x}(s) - \mathbf{y}(t)) \cdot d\mathbf{x} \times d\mathbf{y}}{|\mathbf{x}(s) - \mathbf{y}(t)|^3} \quad (1)$$

In this expression $\mathbf{x}(s)$ are the coordinates of the periodic orbit A , the triple of coordinates being parameterized by the scalar s , and similarly for $\mathbf{y}(t)$ and orbit B .

In three dimensions periodic orbits cannot "pass through" each other. The simple reason is that if they could, two different orbits would at some stage have a point in common. This point, being on two different orbits, would not have a unique future. This violates the determinism property characteristic of all sets of ordinary differential equations used by scientists to model physical processes.

The linking numbers of all pairs of periodic orbits (stable or unstable) in the strange attractor are topological invariants so long as the orbits exist, since the flow was in a three-dimensional phase space $D^2 \times S^1 \subset R^3$. As a result it was possible to use the set of linking numbers for these orbits as a way to characterize/identify the strange attractor of the laser with modulated losses (or at least the mathematical model that described this laser). A table of linking numbers for orbits up to period eight that were extracted from the Belousov-Zhabotinskii strange attractor is provided in Fig. 3

Since the attractor had "a hole in the middle" it was possible to construct another and even more powerful set of topological invariants. Our first impulse was to call these "winding numbers" but Tredicce strongly advised us against using that already appropriated name. Instead, we named them *relative rotation rates*, in recognition of their origin. These are fractions that indicate the average rotation of orbits around each other per topological period.

To be explicit, we can construct a Poincaré section for the flow by hinging a half-plane on an axis and passing the axis through the hole in the attractor. Then each time the trajectory "goes around" it will intersect the half-plane, or Poincaré section, once and from the same side. A period- p orbit A will intersect the Poincaré section p times before repeating itself. Similarly, a different orbit B of minimum period q will have q distinct intersections. All intersections must be distinct, otherwise

Orbit	Symbolics	1	2	3	4	5	6	7	8a	8b
1	1	0	1	1	2	2	2	3	4	3
2	01	1	1	2	3	4	4	5	6	6
3	011	1	2	2	4	5	6	7	8	8
4	0111	2	3	4	5	8	8	11	13	12
5	01 011	2	4	5	8	8	10	13	16	15
6	011 0M1	2	4	6	8	10	9	14	16	16
7	01 01 011	3	5	7	11	13	14	16	21	21
8a	01 01 0111	4	6	8	13	16	16	21	23	24
8b	01 011 011	3	6	8	12	15	16	21	24	21

Fig. 3. Orbits up to period eight were extracted from the experimental strange attractor of the Belousov-Zhabotinskii chemical reaction. The orbits are labeled by a symbol sequence according to where the successive iterations appear on the first return map. (M indicates that iterate occurs at the critical point.) The linking numbers of pairs of orbits are provided in this symmetric table. The linking numbers are orbit pair invariants. Self-linking numbers of the individual orbits are on the diagonal.

determinism is violated: One point on the Poincaré surface of section would have two distinct futures.

The next step is to connect one of the intersections of one orbit with one of the intersections of the other orbit by a directed line segment (an arrow) in the Poincaré section. As time evolves this arrow will move. We can imagine this process happening in the half plane as we sweep the half plane around its axis which goes through the hole in the attractor. If we rotate the half plane through $p \times q$ full rotations, the arrow will come back to its original orientation. This means that it has rotated through an angle of $2\pi n$ radians, or undergone n full rotations in the moving plane. The *average* number of rotations per period is $n/(p \times q)$. This fraction is the relative rotation rate for the given initial conditions on orbits A and B .³³ If $i = 1, 2, \dots, p$ and $j = 1, 2, \dots, q$ specify the intersections of the orbits A and B with the Poincaré section, a relative rotation rate $RRR_{i,j}(A, B)$ can be constructed for each pair (i, j) of initial conditions. There is a simple relation between relative rotation rates and linking numbers:

$$LN(A, B) = \sum_{i=1}^p \sum_{j=1}^q RRR_{i,j}(A, B) \tag{2}$$

Relative rotation rates can even be computed for an orbit with itself. The self-relative rotation rates possessed a very attractive feature: They could be used to distinguish between orbits with nonzero topological entropy and “laminar” orbits with zero topological entropy.³³

This topological index was computed for a number of orbits extracted from the model for a laser with modulated losses. The orbits in this set were identified by a name consisting of a symbol string consisting of 0s and 1s. The symbols were

determined from the location of the intersection of the orbit on the first return map which looked basically like a logistic map. Tables of relative rotation rates are useful complements to tables of linking numbers such as the table in Fig. 3: They provide more information.

Periodic orbits in the Smale horseshoe were also located and identified with a symbol name. The relative rotation rates of these orbits in a suspension of the flow with no twisting (zero global torsion) were also computed.³³ Not all of the orbits predicted by Smale horseshoe dynamics were found in the model of the laser with modulated losses. We compared the table of relative rotation rates (also tables of linking numbers) of the orbits found in the laser model with tables for the appropriate subset of orbits in the Smale horseshoe flow suspension and found, to our delight, complete agreement.³³

With mounting self-confidence we proposed that relative rotation rates could be used as fingerprints to identify strange attractors.³⁴ In fact, this idea was taken one step further: That the integers associated with relative rotation rates (or linking numbers) could be used to *classify* strange attractors.³⁵ We had already identified flows that followed a “Smale horseshoe scenario”,³³ a “Duffing oscillator scenario”,³⁴ and a “van der Pol scenario”, each with distinct sets of linking numbers and relative rotation rates. Why couldn’t these indices be used to classify/distinguish one chaotic attractor from another? It seemed reasonable to hope so.

4. Branched Manifolds

At this point my first hope (it was mine: my colleagues Solari and Mindlin weren’t nearly as loopy as me) was that we could create a dictionary of scenarios (or processes, or mechanisms) and for each construct a table of linking numbers and/or relative rotation rates for lots of low-period orbits. Then when confronted with experimental data we could pull out the unstable periodic orbits from the data, compute their topological invariants, and then thumb through the dictionary comparing tables until we found an agreement. This program represented a lot of work.

While we were mulling over implementing this program a better solution became available. The beautiful work of Joan Birman and Robert Williams^{36,37} gradually seeped into our consciousness. The time scales for implementing the “dictionary program” and for understanding the Birman-Williams Theorem were comparable, but the level of rewards for the latter far outweighed the former. The net result was that we understood the Birman-Williams Theorem at a level sufficient to apply it to our physical needs (“experimentalists’ level”).

What we understood is this. Suppose we have a dissipative flow in three dimensions whose trajectory “is a strange attractor.” There is one positive Lyapunov exponent $\lambda_1 > 0$, one negative Lyapunov exponent $\lambda_3 < 0$, and one zero exponent $\lambda_2 = 0$ “along the direction of the flow”. The dissipative nature of the flow requires $\lambda_1 + \lambda_2 + \lambda_3 < 0$. Then it is possible to project points in the phase space “down”

along the direction of the stable manifold. This is done by identifying all points with the same future:

$$x \simeq y \quad \text{if} \quad \lim_{t \rightarrow \infty} |x(t) - y(t)| = 0 \quad (3)$$

In this expression $x(t)$ is the future in the phase space of the point $x = x(0)$ under the flow. This Birman-Williams identification effectively projects the three-dimensional phase space down to a two-dimensional set that is a manifold almost everywhere. The two dimensions that remain correspond to the flow direction (with $\lambda_2 = 0$) and the unstable or stretching direction (with $\lambda_1 > 0$). The “almost everywhere” refers to zero- and one-dimensional sets where the “manifold condition” is violated. The points in the projection describe where the flow “splits” and the branch lines describe where flows from two distinct parts of phase space are “squeezed together”. These rigorous mathematical structures were prefigured at an intuitive level by Lorenz²⁷ and Rössler²⁸ long ago. The Rössler attractor and its branched manifold are shown in Fig. 4 and the Lorenz attractor and its branched manifold are shown in Fig. 5.

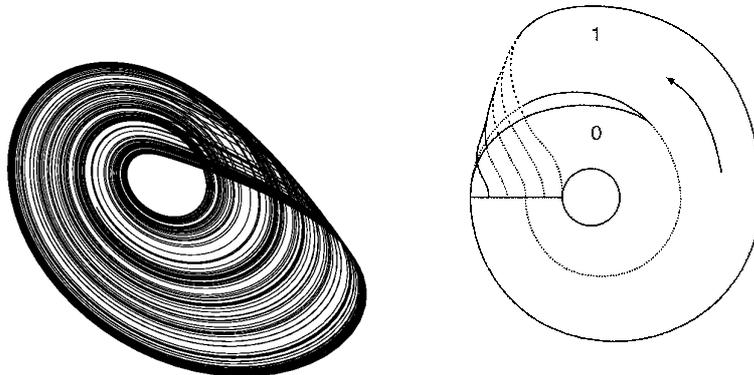


Fig. 4. The Rössler attractor is shown on the left in its projection onto the x - y plane. Its branched manifold is shown on the right. The two branches are labeled 0 and 1. The integer indicates the torsion of the branch. The two branches split at the “splitting point” (near the arrowhead) and join at the branch line.

Branched manifolds are useful constructions for distinguishing among different mechanisms that generate strange attractors. Four of the most studied strange attractors are those associated with the Lorenz, Rössler, Duffing, and van der Pol dynamical systems. The branched manifolds that describe these strange attractors are shown in Fig. 6.⁴ These four branched manifolds are topologically inequivalent. “Equivalence” is by isotopy: Two things are isotopic if it is possible to mold one into the other without tearing or gluing it. As a result, identifying the branched

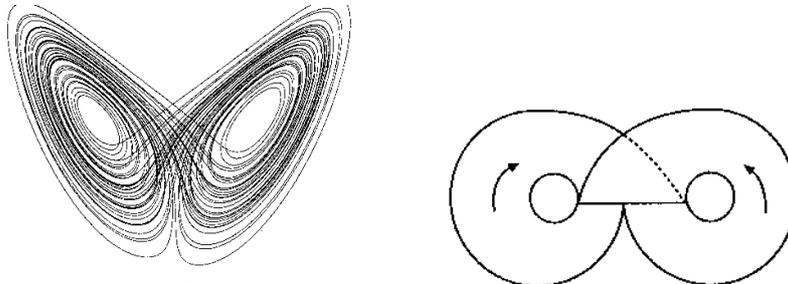


Fig. 5. The Lorenz attractor is shown on the left in its projection onto the x - z plane. Its branched manifold is shown on the right. Neither branch exhibits a twist in this representation of the branched manifold. The two branches split at the “splitting point” which is shown “in” the branch line.

manifold that describes a strange attractor is a powerful tool for distinguishing one (class of) strange attractors from another.

We should point out that branched manifolds can be constructed from ‘stretching’ and ‘squeezing’ units. These units are shown in Fig. 7. There are two simple rules for this aufbau construction:

1. Outputs to inputs;
2. No free ends.

An enjoyable way to construct elegant branched manifolds is to dump a bushel of stretch/squeeze units in front of an enthusiastic class of kindergarteners, along with the instructions above. There is no guarantee that scientists can find a physical system described by the resulting artworks.

There are technical aspects in the statement of the Birman-Williams Theorem that we chose to ignore. There are three in fact. For the proof of the theorem, the flow is assumed to be:

1. Hyperbolic
2. In R^3
3. Dissipative.

We found it useful to ignore the hypothesis on which the theorem is based (at our own peril) for the following reasons:

Hyperbolic: In physics, both in theory and experiment, we have never seen a hyperbolic attractor. Every experimental chaotic attractor that we have seen is continually undergoing bifurcations as the external controls are varied. For example, the logistic map $x' = \lambda x(1 - x)$ is hyperbolic only for $\lambda > 4$. It is stable (dynamically but not structurally) only at the knife edge $\lambda = 4$ and a strange *repellor* for $\lambda > 4$. If we were to insist on the assumptions undergirding the theorem it would not be useful to us as physicists.

In R^3 : Three dimensional models (e.g., Lorenz²⁷ and Rössler²⁸) generate three

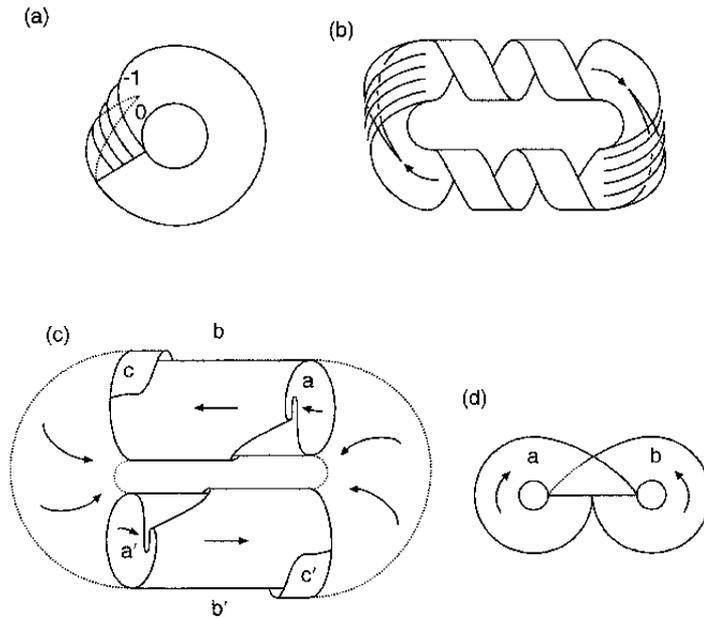


Fig. 6. Branched manifolds for: (a) Rössler; (b) Duffing; (c) van der Pol, and: (d) Lorenz attractors. It is clear that no topological move (stretching and folding are OK; tearing and gluing are not) can transform any one of these to any of the others. Therefore it is not possible to find a 1:1 coordinate transformation that converts any of these dynamical systems to any of the others.

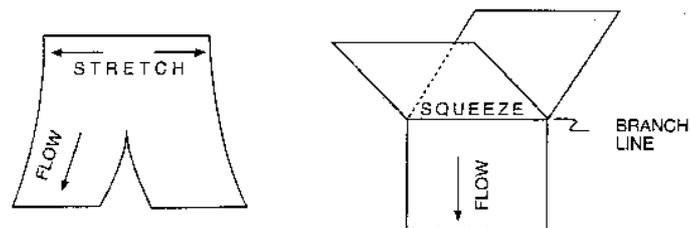


Fig. 7. Stretching and squeezing units serve as the basic building blocks for constructing branched manifolds.

dimensional data sets, or scalar data sets (e.g., $x(t)$) that can be embedded in three-dimensional spaces. But physical processes are often described by high dimensional sets of ordinary differential equations, or even by partial differential equations. In short, having collected a scalar data set there is no

guarantee at all that it is generated by a three-dimensional set of equations. Even so, if we can find a three dimensional embedding the gods have smiled at us, and there is every anticipation of further smiles allowing the theorem to “work” even though the origin of the data set does not conform to the assumptions underlying the proof of the theorem.

Dissipative: One of the motivating ideas for the theorem comes from classical physics in a conservative manifestation. The magnetic field lines surrounding an infinitely long straight wire carrying a uniform current are closed loops described by two continuous parameters: radius and position along the axial direction of the wire. When the wire is bent most of the closed loops break: typically only a countable measure-zero set remains closed. This is true in particular when the wire is tied into a figure eight knot (carrying a supercurrent). The unbroken magnetic field lines are rigidly organized among themselves. The organization can be discerned by the linking numbers that pairs of closed loops exhibit. Further, these closed magnetic field lines can be isotoped down to the figure eight branched manifold without in any way altering these linking numbers.^{4,36,37} In this conservative case the thrust of the Birman-Williams Theorem is valid, even though the proof doesn’t cover this case.

I think the Birman-Williams Theorem is more powerful than it appears. Two of the three conditions discussed above are no longer an impediment to important physical applications.

Hyperbolicity: We avoid this constraint by falling back on a pruning argument. This will be explained in Section 7.

In R^3 : The theorem is true for *strongly attracting* dynamical systems with one positive Lyapunov exponent. What does this mean? If the spectrum of Lyapunov exponents for a flow in R^n satisfies $\lambda_1 > \lambda_2 = 0 > \lambda_3 \geq \lambda_4 \cdots \geq \lambda_n$, the system is strongly attracting when $\lambda_1 + \lambda_2 + \lambda_3 < 0$.³⁸

As for the third (dissipativity), this is needed for the Birman-Williams projection (Eq. (3)) to work. However, it may not be necessary to project the flow to a branched manifold in order to describe the organization of the unstable periodic orbits by the properties of a branched manifold. Indeed, the original inspiration of the unbroken magnetic field lines and the figure-eight branched manifold shows that there is at least one case where the theorem “works” despite the fact that the underlying assumptions are not satisfied. I suspect there are many more such cases (one for every oriented knot); perhaps even a theorem.

5. Topological Analysis Program

Once we understood what information we wanted to extract from the data, it was time to act. We acted first with data from Tredicce's laser. Of course, analyzing experimental stuff is not the same as analyzing computer generated data. Our attempts to extract a topological understanding of the laser data failed. In retrospect the reason was easy to see in the data. The laser acted somewhat like a relaxation oscillator. A large percentage of the time ($\simeq 40\%$) the intensity output (this is the observed variable) was very low. At such times the intensity was recorded in the lowest channel of a multichannel analyzer. When it came time to determine the topological organization of the unstable periodic orbits extracted from the data, we found many crossings occurred in the "lowest channel". Result: we were unable to determine accurately or honestly the crossing properties in low intensity regions. A request to redo the experiment was met with amused incredulity. It had long since been dismantled and pieces sent to new experiments, returned to previous owners, broken, or otherwise no longer available - a typical situation in a working experimental laboratory.

Plan B involved placing an APB for experimental data at each of the Nonlinear Dynamics meetings that we attended during this period. Lathrop and Kostelich responded to this plea. They had used periodic orbits in order to characterize an experimental strange attractor for the Belousov-Zhabotinskii reaction,³¹ and they provided us with a sample of these data. The data had been taken by Swinney's group in Texas.⁹⁻¹¹

The data were inordinately clean. Even so, we had a few problems attempting to make our favored embedding. Ultimately, these problems were resolved.³⁹

We took the following steps to create the first topological analysis of experimental data.

- (1) A set of unstable periodic orbits was extracted directly from the scalar time series *before* an embedding was created. This process is indicated in Fig. 8.
- (2) A suitable embedding was found. We used an integral-differential embedding. In this embedding each orbit was identified by a symbol string obtained from a suitable return map. Only two symbols (0, 1) were sufficient since the return map had two branches.
- (3) A table of linking numbers was made. This involved all orbits found in the first step, and the linking number was determined in the three-dimensional space constructed in the second step.
- (4) A branched manifold was identified that "explained" these linking numbers. Our branched manifold had two branches since only two symbols were required to represent the trajectory described by the data set.

A gloss on the last step is appropriate at this point. The branched manifold supported more orbits of any period p than we actually extracted from the data.

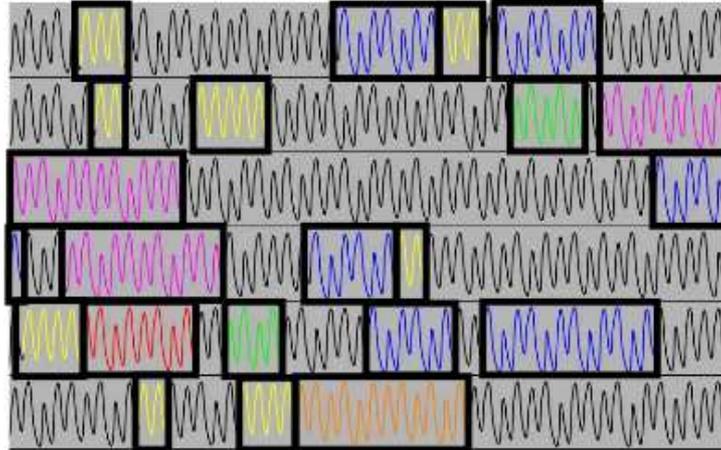


Fig. 8. (color online) These data were taken by Lefranc, Hennequin, and Glorieux,⁴¹ who redid Tredicce's laser experiment using a logarithmic amplifier-detector to climb out of the lowest-channel bind. The different-period orbits are coded with different colors. The identification precedes the embedding step. These authors also showed that the fractal dimension computed with and without the data processing differed substantially, 'contradicting' a theorem that fractal dimension is an invariant.

We regarded this as a case of the absent orbits having been "pruned" from the hyperbolic limit. We were able to make a 1:1 correspondence between the orbits extracted from the data and an appropriate *subset* of orbits on the branched manifold. The correspondence was by symbol name. A table of linking numbers was constructed for the experimental orbits. This was done in two ways. Two orbits were superposed, the signed number of crossings counted, and the result divided by two.³³ This counting-of-integers method was double checked by computing the Gauss linking integral (Eq(1)). In the three-dimensional embedding each orbit was defined by a vector $(x(s), y(s), z(s))$ of coordinates, where s is a useful parameter. A code for computing the Gauss double integral was developed, extensively tested, and then applied to the embedded experimental data.

On the branched manifold side things took a different turn. We began by determining where each periodic orbit would fall on the branched manifold. This is straightforward but tedious. We continued by superposing the linear crossing segments and counting signed crossings. This is also straightforward but even more tedious up to period $\simeq 4$, after which it became error prone and almost impossible due to blurry eyeballs. In desperation we resorted to doing this by computer (think Geiger counter!). A code was written to do these things faster and better than a human could. The code did two things. It first located periodic orbits on a branched manifold (applied kneading theory⁴⁰). Then it counted crossings of orbit pairs. Two inputs were required. One was the symbol names of the orbits under

consideration. The second was a quantitative description of the branched manifold. Part of this description was a square $n \times n$ matrix, where n is the number of distinct symbols required to describe a trajectory in the experimental data as well as the number of different branches in the branched manifold ($n = 2$ in our case). The second was an $n \times 1$ array defining how the two branches are joined at the branch line.^{4,42}

We later realized that these quantitative indices — square matrix and array — initially introduced as an *aide comptoir* in our computer code, actually served as an integer description of the branched manifold, which is itself a mathematically rigorous characterization of the experimental chaotic attractor (up to pruning).

Lastly, we realized that the branched manifold could be identified using only a very small number of unstable periodic orbits extracted from the data. This branched manifold could then be used to *predict* the linking numbers or relative rotation rates of all other orbits and orbit pairs extracted from the experimental data. These predictions could then be compared with those constructed from the experimental orbits. Either there was complete agreement, and we could conclude our identification was correct, or else there wasn't — that is, one or more integers in the linking number tables were not the same — in which case we had to *reject* the hypothesis that our identifications, of orbit labels and/or branched manifold characteristics, was correct. The paper in which we reported the results of our analysis of the Belousov-Zhabotinskii data is also the paper in which we announced the Topological Analysis Program.³⁹

At last there was a rejection criterion for the analysis of chaotic dynamical systems that wasn't subjective. No error bars are associated with integers!

There were additional useful benefits from this program (Principal of Unexpected Beneficial Consequences). Noise usually degrades the analysis of data. In the case of topological analyses, noise makes it more difficult to extract the longer periodic orbits. The most important orbits for this analysis are the lowest period orbits. The effect of decreasing S/N is to reduce the number of *surplus* orbits, so that the rejection step is carried out with fewer orbits. In spite of this, even with a moderate amount of noise there are more than enough orbits to make a branched manifold identification and then to carry out the rejection tests using the remaining orbits.⁴²

Metric analysis methods call for very long, clean, stationary data sets. Any kind of nonstationarity will destroy fractal dimension estimates and seriously degrade dynamical estimates (e.g., Lyapunov exponent spectra). However, it is possible to identify orbits with positive topological entropy in highly nonstationary data. Such identification provides a clear statement that the underlying dynamics is chaotic. Identifications of this nature have been carried out by Amon and Lefranc.⁴³

6. Chess Pieces on the Board

It was at this point that Prof. Ennio Arimondo sent us a gift in the form of Francesco Papoff, laded with data. The data had been taken in Arimondo's laboratory using lasers with saturable absorbers.⁴⁴ Several different absorbers had been used, and for each saturable absorber the laser had been run under a large number (6 - 10) of operating conditions. This provided us with a serendipitous chance to test another of our favorite hunches.

The idea was this. Suppose you have a physical system operating in a chaotic regime. A strange attractor is produced. The attractor can be investigated as described above, yielding a spectrum of unstable periodic orbits and an underlying branched manifold. This latest hunch was that as 'perturbations' are made, for example changing the operating conditions, the underlying branched manifold remains unchanged but the perturbations "push the flow around" on the branched manifold. We hoped to find that *all* data sets, when analyzed, yielded the same underlying branched manifold but the spectrum of orbits extracted from the various data sets changed from one set to the next. Indeed, this is exactly what we found.⁴⁵ Further, the underlying branched manifold was exactly the same as we had previously found for data from the Belousov-Zhabotinskii chemical reaction.

This result, based on the analysis of experimental data, was important in resolving a somewhat philosophical debate on how branched manifolds should be used in physics. This debate is summarized in Fig. 9. On the left in this figure is the Lorenz attractor — its branched manifold is apparent. On the right is a different attractor, the Shimizu-Morioka attractor.⁴⁶ Its branched manifold can also be readily inferred. By unwinding the loops on the left and right it bears a close resemblance to the Lorenz branched manifold, but the unwound (writhing) loops now each have a full twist and describe orbit segments of topological period two. Is it more useful to use this as the branched manifold for the Shimizu-Morioka attractor, or rather regard the periodic orbits in this attractor as a subset of the orbits in the Lorenz branched manifold?

Some argued that the branched manifold used to describe a physical system should be the one that contains only and exactly the periodic orbits in the attractor. The implication is that as operating conditions change and the spectrum of orbits in the attractor varies, the associated branched manifold would undergo mind-boggling contortions. Contortions so complicated as to curtail the use of branched manifolds as a nifty tool for understanding chaotic attractors.

On the opposite side of this opinion was our feeling that a branched manifold should be introduced with one branch for each of the symbols required for a unique description of an arbitrarily long trajectory — or at least of the length measured in an experiment. Such a branched manifold would "contain" all the orbits seen in an experiment as well as a lot more. It would be more useful, we maintained, to regard the missing orbits as having been pruned from the original allowed spectrum

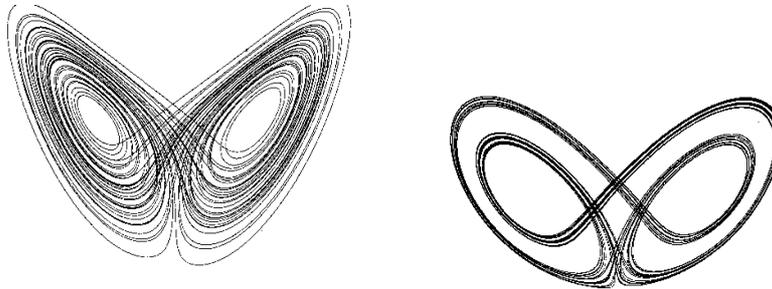


Fig. 9. The Lorenz attractor²⁷ and the Shimizu-Morioka attractor.⁴⁶ It can be argued that the two branched manifolds are different, or that the periodic orbits in the Shimizu-Morioka attractor are simply a subset of those in the Lorenz attractor. In the latter case the Lorenz branched manifold can be used to compute the topological indices of those orbits in the Shimizu-Morioka attractor that are not pruned from the Lorenz attractor.

(technically, a “full shift on n branches”⁴⁷). With this assumption, all the remaining orbits are organized in exactly the same way as in the branched manifold with unpruned spectrum. And perestroikas — experimentalists love to tweak knobs connected to control parameters — would generally leave the branched manifold untouched while changing the spectrum of orbits in the attractor. This is the source of our idea that changing control parameters serves to push the flow around on an unchanged underlying branched manifold. This simple view is very useful.

7. Forcing

While extracting periodic orbit surrogates from the 25 data sets that Papoff brought to us, Mindlin, Papoff, and Ricardo Lopez-Ruiz noticed that orbits had a social life of their own. In particular, they observed that when one particular orbit was present it was invariably accompanied by a handful of other orbits — always the same handful. Several different orbits possessed this property. The presence of some orbits seemed to “force” the presence of other orbits. This observation cried out for a more careful look.

We approached this problem from a topological perspective, since topology had already been so kind to us. When orbits are created, they are created in saddle-node pairs (neglect period-doubling for the moment). We computed the linking numbers of the pair A_n, A_s (A_n is the node of the saddle node pair of orbits A) with the pair B_n, B_s and arranged the results in a 2×2 matrix (we like matrices):

$$\begin{array}{c|cc} & B_n & B_s \\ \hline A_n & L(A_n, B_n) & L(A_n, B_s) \\ A_s & L(A_s, B_n) & L(A_s, B_s) \end{array} \quad (4)$$

Depending on the equality or inequality of the four integers in this matrix, it was

possible to determine if orbit B could exist before orbit A had been created, or vice versa (see Fig. 10). In this way we were able to piece together an orbit forcing diagram.⁴⁸ We carried the calculations out up to period eight. We were also able to develop the idea of a useful subset of orbits, which we called, because of our background, a *basis set of orbits*.

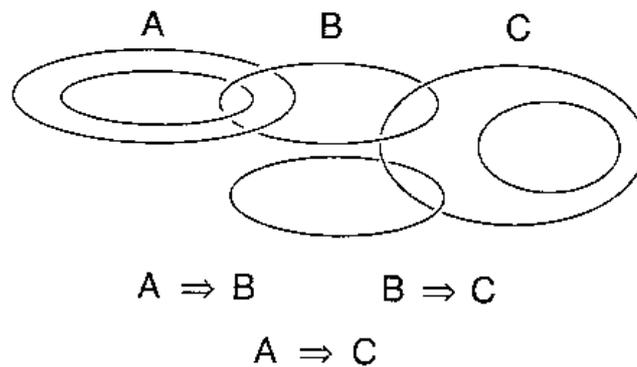


Fig. 10. Orbit pair B must be created before orbit pair A . The two orbits A_n and A_s have different linking numbers with the orbits B_n and B_s . Similarly, orbit pair C must be created before orbit pair B . The existence of the pair A forces the presence of pair B , and $B \Rightarrow C$.

The idea is as follows. Organize all the periodic orbits extracted from an experimental data set. Organize them according to their two-dimensional entropy, using the one-dimensional entropy as a tie-breaker.^{4,42,48} For example, we found these orbits in one set of experimental data:

$$\underline{2_1}, \underline{4_1}, \underline{8_1}, \underline{6_1}, \underline{8_2}, \underline{7_1}, \underline{5_1}, \underline{8_3}, \underline{3_1}, \underline{6_2}, 6_3, 7_5, 4_2, 8_9, \underline{8_4}, 8_7, 7_6, \underline{7_2}, \underline{7_3}, 7_4, \underline{8_5}, 8_6, 8_8, 5_2$$

These orbits are identified by their order of creation in the logistic map. Then, starting with the “highest” orbit (5_2) with the largest entropies, we remove that orbit and “kill off” all the orbits that it forces (these are underlined). If there are any orbits left (yes unless the two-dimensional entropy is equal to the one-dimensional entropy), we continue the process:

$$\underline{6_3}, 7_5, 4_2, 8_9, 8_7, 7_6, 7_4, 8_6, 8_8$$

And again, until no orbits are left. Then the small set of orbits that have been removed (as opposed to “killed off”) consists of the basis set of orbits on the branched manifold that describes the chaotic attractor. These are the orbits

$$8_7, 7_6, 7_4, 8_6, 8_8, 5_2 \tag{5}$$

In truth, this argument works up to whatever finite period forcing information is available.

A lower bound on the topological entropy of the flow can be obtained by computing the topological entropy of the braid containing the basis set of orbits.

‘Forcing’ is a very difficult problem, and in truth our approach is probably the least effective that has been found to work. Other approaches^{49–54} are more effective, but much more complicated. The problem of forcing, even on the two branch manifold describing Smale horseshoe dynamics, is still open.

8. Branched Manifold Perestroikas

Normally, small perturbations produce changes in the spectrum of orbits that a branched manifold can support. But sometimes perturbations are not small. Under these conditions the branched manifold will change. What this amounts to is that more symbols are required to uniquely label a chaotic trajectory. Or perhaps fewer symbols are required. Yet more generally, the spectrum of symbols requires changes. Since symbols correspond to branches, such perturbations lead to changes in the structure of the branched manifold underlying the description of a chaotic attractor and its perestroikas.

Such changes were studied extensively for the periodically driven Duffing oscillator as a function of changing driving frequency.^{55,56} The modifications that occur in the dynamics, the spectrum of stable and unstable periodic orbits, and the number, nature, and organization of the branched manifolds that describe the attractor, and the topological indices that are observable occur with regular predictability. In short, each time the external driving frequency passes through a $1/n$ subharmonic of the natural frequency of the undriven nonlinear oscillator, the underlying branched manifold undergoes another full twist. The global torsion changes systematically by ± 1 in each subharmonic window.

It seems almost as if there is a “branched manifold” that assumes the form of an enormous scroll with branches labelled by successive integers $0, 1, 2, \dots$, and the flow is constrained to travel in only a small number of adjacent branches for any value of the control parameters. As the controls are changed the flow is directed over a limited number of contiguous branches, for example $(0, 1) \rightarrow (0, 1, 2) \rightarrow (1, 2) \rightarrow (1, 2, 3) \rightarrow (1, 2, 3, 4) \rightarrow (2, 3, 4) \dots$. This systematic behavior is a consequence of continuity. This behavior has variously been called the “jelly roll scroll” (in the US) and the “gateau roulé” (in France). One important question is whether the scroll rolls up from outside to inside (yes in the cases we have investigated) or from inside to outside. These two cases can be distinguished by computing linking numbers of appropriate orbits.^{4,42,55,56}

Shortly after this study, scrolling behavior was observed in experiments performed on a periodically driven Nd-doped fiber optic laser⁵⁸ As the external drive frequency descended through the subharmonics of the natural resonance the branched manifold describing the chaotic behavior became more and more wound up. The systematics of this behavior is indicated in Fig. 11. This winding-up

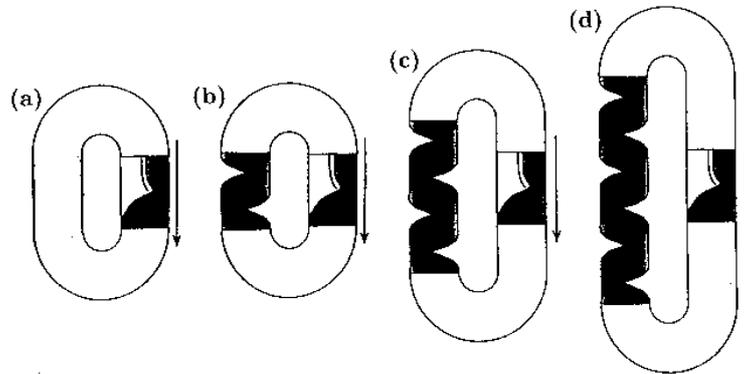


Fig. 11. As the ratio between the natural frequency (period) of an oscillator and the frequency (period) of an external driver decreases (increases), the strange attractor that may exist becomes increasingly wound up. The global torsion increases systematically as the ratio $T_{natural}/T_{drive}$ increases through integer values.

phenomenon is a typical feature when there is a competition between some natural resonance frequency and an externally imposed driving frequency.⁵⁷ It has been seen in the analysis of chaotic signals from dogfish, catfish, and paddlefish^{59,60} and many laser experiments.^{61,62}

Instead of winding up in a *gateau roulé* structure, it is possible that three continuous branches of a branched manifold could fold in an *S*-shaped structure. This has finally been found in one of the subharmonic domains studied by Javier Used and Juan Carlos Martin in their periodically driven Erbium doped fiber optic laser.⁶² This work is summarized in another contribution to these Proceedings.

9. Branched Manifolds Describe Mechanism

Branched manifolds are exactly the right tool for describing the *mechanism* that acts to create a strange attractor and at the same time to organize all the unstable periodic orbits in it. By mechanism we mean the stretching and folding and/or the tearing and squeezing processes that occur repetitively in the phase space.^{4,42}

We illustrate two different mechanisms in Figs. 12 (Rössler mechanism) and 13 (Lorenz mechanism). In Fig. 12 a blob of points in the phase space at (a) is stretched along one direction (with $\lambda_1 > 0$) and flattened in another (with $\lambda_3 < 0$) while being displaced in a third direction (with $\lambda_2 = 0$). When the flow exists in a bounded region of the phase space, some mechanism must exist to return this set of points to its initial neighborhood. One mechanism involves a simple fold, shown beginning at (c) in this figure. Eventually the folded-over set of points returns to its initial neighborhood as (c) \rightarrow (d) \rightarrow (a). This process repeats indefinitely, building up a flakey structure in the squeezing (λ_3 direction) known as a fractal. In the case

depected in Fig. 12 the mechanism is a simple stretch and fold. This is represented by a simple branched manifold with two branches.

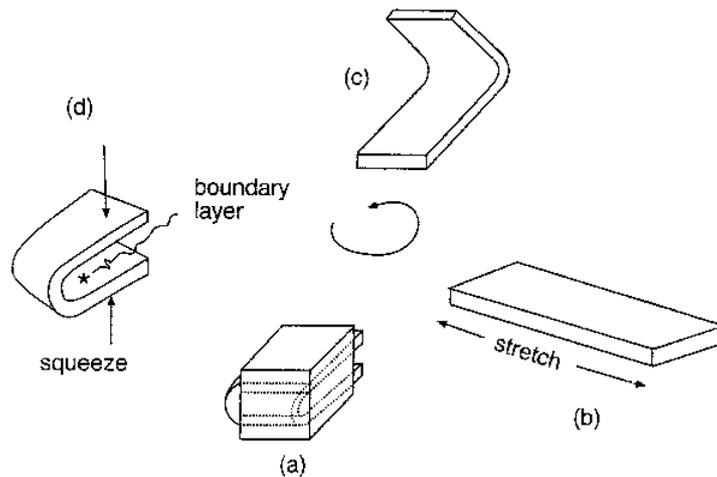


Fig. 12. A set of initial conditions at (a) is deformed by stretching in one direction and squashing in another. As the flow progresses, folding begins to occur at (c) and continues on through (d). This deformed set of initial conditions finally returns to its initial neighborhood (a), where the processes is repeated *ad infinitum*.

Another mechanism is illustrated in Fig. 13. In this case a set of initial conditions (the cubes) begin to flow but run into a “buzz saw” that cuts the set into two or more pieces. These pieces move off into different directions in the phase space, where they encounter and are squeezed together into other blobs of points, some with different initial conditions. Eventually these sets of points encounter their initial neighborhoods, and the process continues over and over again.

Branched manifolds summarize in a simple and visual way the mechanisms that exist and occur repetitively to build up each strange attractor.^{63,64}

10. Bounding Tori

Branched manifolds place constraints on the periodic orbits that can be created or destroyed during a perestroika. At some point we began to wonder if there were larger structures that placed analogous constraints on branched manifolds themselves during a perestroika.

In fact, there are. We began by asking about the topological properties of a “smoothed out” version of a messy fractal chaotic attractor. We could smooth a strange attractor by surrounding each point in the attractor by a ball of ϵ radius (ϵ is “small enough”) and constructing the union of all such three-dimensional balls.^{65,66}

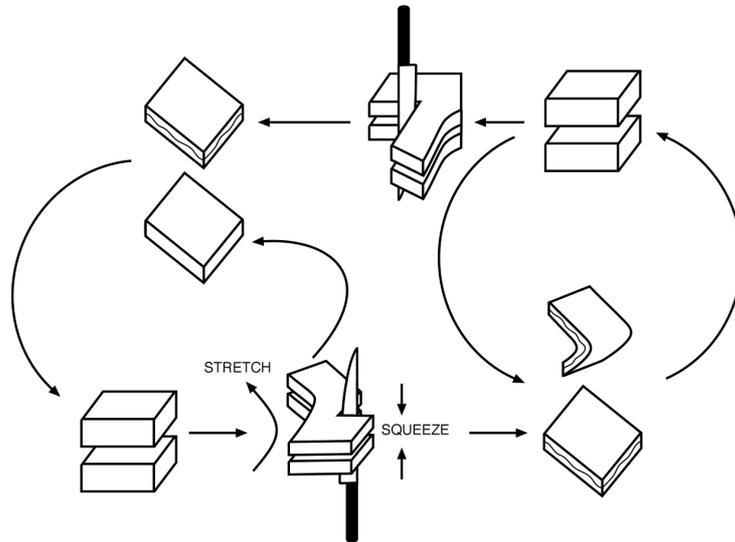


Fig. 13. Sets of initial conditions (cubes) are “sliced”, by running into an axis with a stable and an unstable direction (the z -axis for Lorenz-like systems), for example. The different parts flow off in different directions in the phase space, where they may encounter other sliced parts from different regions of phase space. These are squeezed together and eventually return to regions they originated from (recursion).

When done correctly, the union was a three dimensional manifold. Since we were looking for some structure to enclose or surround the attractor and its branched manifold, we were naturally led to consider the boundary of this manifold. Once again, we wound up talking to topologists. These boundaries were two-dimensional, orientable, and bounded — therefore uniquely tori. A torus is shown in Fig. 14 along with some of the important closed loops on it. A topological torus is characterized by one number, its genus g . Our torus surfaces were “dressed” with the flow that generated the chaotic attractor inside the torus, restricted to the surface. The immediate result was that a strange attractor could be described by an integer, g , the genus of the torus surrounding it, together with another more complicated discrete index that describes the flow on the surface. The second index is not an integer but a transition matrix related to the permutation group P_{g-1} .⁶⁵⁻⁶⁷

On the surface the flow has some stagnant points, or fixed points. All fixed points on the surface arise when the tangent vector (recall, $\lambda_2 = 0$) of the flow generating the attractor is perpendicular to the surface. At such points the stability of the restricted vector field is governed by the two remaining eigenvalues, $\lambda_1 > 0$ and $\lambda_3 < 0$. As a result, the index of each fixed point on the surface is -1 and the sum of these indices is related to the genus of the torus by $\sum_{\text{all f.p.}} (-1)^1 = -2(g - 1)$. A lot of elegant topology due to Euler and Poincaré goes into this result.⁶⁸

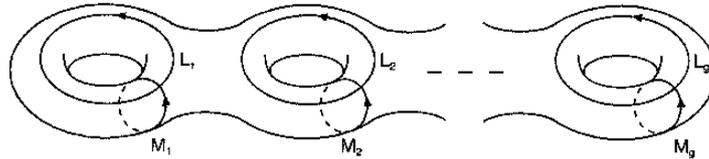


Fig. 14. A torus in R^3 is completely identified by the number of holes in it. This integer index is called its genus. A normal tire tube is a genus-one torus and a sphere has genus zero. A torus of genus g has two useful sets of closed loops on it. These form the generators of its homotopy group. The meridians (M_1, M_2, \dots, M_g) can be chosen to bound a disk that lies entirely within the torus. The longitudes (L_1, L_2, \dots, L_g) bound disks that lie entirely outside the torus in the case shown. This isn't always true: to see how this can fail, see Fig. 17. It is always possible to choose g meridians and g longitudes so that they are independent, meridians have no intersections with other meridians and similarly for longitudes, and each meridian intersects only one longitude at one point. The tori bounding strange attractors are “dressed” on their surfaces with the flow that generates the attractor.

Working by analogy again, we asked if there is an “aufbau principal” for bounding tori the way there is one for branched manifolds. In the latter case a branched manifold can be built up by joining together stretching and squeezing units (c.f., Fig. 7). We didn't have to look too far (not beyond string theory) to find what we needed. The building units we needed were called variously “pairs of pants” or “trinions”. These come in two varieties, as do the basic building blocks of branched manifolds. The two building blocks are shown in Fig. 15. Stretching trinions have one input port and two output ports. These contain the branched manifold stretching units (c.f., Fig. 7). Joining trinions have two input ports and one output port. These contain the branched manifold squeezing units.

We found that it is possible to build up any genus- g torus ($g > 1$) using $g - 1$ pairs of these units. Each pair contains one stretching unit and one joining unit. These units must be connected together following the usual rules:

- (1) Outputs connect to inputs.
- (2) No free ends.

Two pairs of stretching/squeezing units can be used to build up the genus-three torus that bounds the Lorenz attractor, as shown in Fig. 16

An added benefit of this decomposition is that it is now possible to describe the global Poincaré surface of section for a low-dimensional chaotic attractor. Determine its bounding torus. Partition it into $g - 1$ pairs of stretching/joining units. The output port of each joining unit (equivalently, the input port of each splitting unit) is a disk that is one component of the Poincaré surface of section. The section itself is the union of these $g - 1$ disks. A topological period (as opposed to a temporal period, measured in seconds) is a transition from the section to itself: this means from any one of the disks to the next, under the flow. The Poincaré section of the

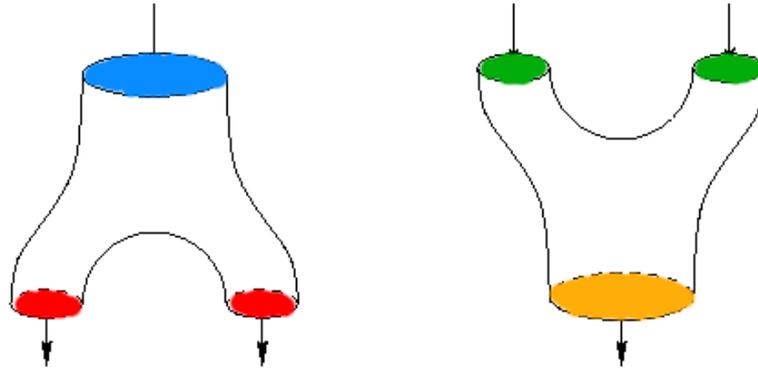


Fig. 15. (color online) The basic building blocks of bounding tori are stretching (left) and squeezing (right) trinions. A genus- g torus is built up via the aufbau principal using $g - 1$ pairs of stretching/squeezing trinions, $g > 1$. The construction is: Inputs to outputs; with no free ends at the end of the construction. Each join is colorless. The global Poincaré surface of section is the union of $g - 1$ disks. The disks can be taken as the input ports of the splitting trinions or the output ports of the joining trinions.

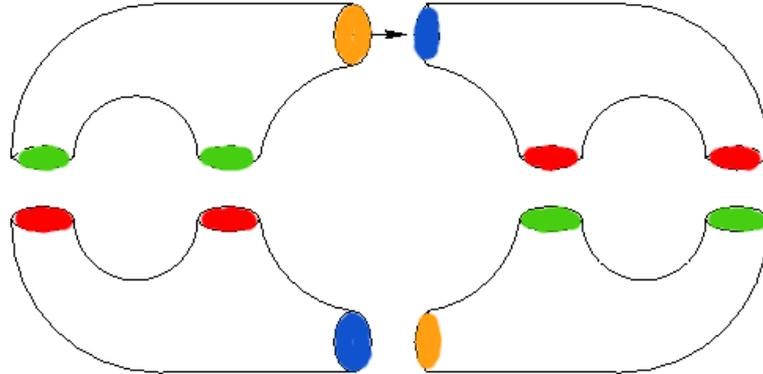


Fig. 16. Two pairs of stretching and squeezing trinions are used to build up the genus-three torus that encloses the Lorenz attractor. Output ports of squeezing trinions flow to the input ports of stretching trinions, etc. All connections are colorless. The global Poincaré surface of section is the union of disks. The disks can be taken as either the input ports of the stretching trinions or the output ports of joining trinions.

Lorenz attractor is the union of two disks, as seen in Fig. 16.

All of the experimental attractors that we have analyzed that have a “hole in the middle” live in genus-one tori. The mathematical attractors with this feature include the Rössler, Duffing, and van der Pol attractors. The latter is bounded by *two* genus-one tori, one outside and one inside the attractor. The Lorenz attractor lives inside a genus-three torus. A number attractors studied by Aziz-Alaoui⁶⁹ live in higher genus tori.

The closed magnetic field lines surrounding a knot tied into the shape of a figure eight knot have the topological organization defined by the figure eight branched manifold. This lives inside a bounding torus of genus $g = 9$. It seems an interesting idea to relate closed magnetic field lines generated by supercurrents in wires tied into the form of various knots with the associated knotholder and relate the genus of the knotholder with the original knot.

We point out that tori can be embedded into R^3 in a multitude of bizarre ways, as indicated in Fig. 17.

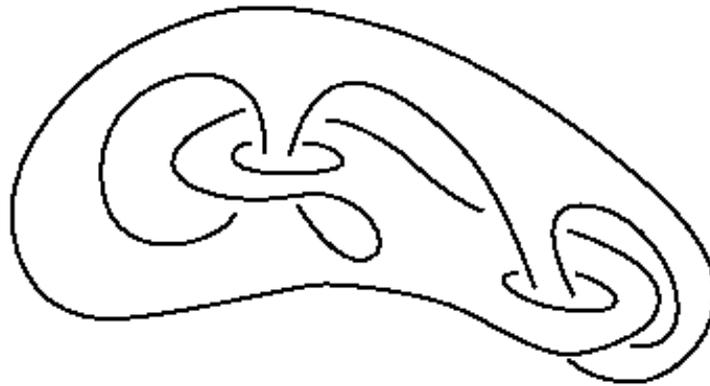


Fig. 17. Tori may be embedded into R^3 in a variety of different ways. For this genus-four torus all meridians bound disks that lie entirely within the torus. However, each of the four simplest longitudes links one other longitude.

11. Four Levels of Structure

We now can address the levels of structure available to describe a strange attractor. We regard attractors that can be deformed into each other smoothly — no cutting or tearing, no creating or annihilating orbits — as isotopic or equivalent. We ask: how do we distinguish (topologically) inequivalent attractors.

At present, there are four levels of structure. Each level is discrete. We describe these levels in some semblance of order.

At the lowest level there is a basis set of orbits (to any finite period). This set determines all orbits that are present in the attractor, up to some period.

At the next level are branched manifolds. These serve as a rigorous caricature for the strange attractor. Branched manifolds constrain the evolution of basis sets of orbits under perestroikas.

Above branched manifolds in this hierarchy are bounding tori. Bounding tori constrain the evolution of branched manifolds under perestroikas.

Bounding tori live in R^3 . A torus ($g = 1$) can be embedded in R^3 in many inequivalent ways. To determine how many, imagine shrinking the torus surface down to an oriented curve in R^3 . Then there are as many ways to embed a torus into R^3 as there are oriented knots in R^3 . Each such knot defines a torus and within each such torus we can construct a strange attractor. We can do better: we can embed a $g = 1$ strange attractor in each knotted torus in such a way that all the embedded strange attractors are diffeomorphic. They are not isotopic, since the knots cannot be deformed into each other.⁷⁰

For a genus one torus containing a strange attractor there are as many embeddings as there are oriented knot types. Although we cannot yet distinguish inequivalent knots by any known index, the number of knot types is discrete.

Similar arguments apply to tori of genus- g . Each genus- g bounding torus can be embedded in many different ways in R^3 (c.f., Fig. 17). This discrete index we call, for want of a better name, the generalized knot index. Constructing a generalized knot index is an even more difficult problem than constructing a knot index for the simpler case of the genus-one torus.⁷¹

12. Symmetry

We were always on the lookout for new strange attractors. So when Christophe Letellier asked about the relation between symmetry and chaos an immediate bond was formed between us. The first question we asked is: “How does Cartan’s theorem about the relation between covering groups and their images play out in dynamical systems theory?”⁷² Of course, at that time we had no clue..., and for this reason the chase was exciting.

We were first motivated by the way we could identify symmetric points in the Lorenz attractor and make it “look like” the Rössler attractor. Eventually these visceral feelings gave ground to a more quantitative approach.⁷³⁻⁷⁵

In one direction (image direction) things are relatively simple. The Lorenz equations are unchanged (equivariant) under the transformation $(X, Y, Z) \rightarrow (-X, -Y, +Z)$. By identifying a point with coordinates $(-X, -Y, Z)$ with its partner (X, Y, Z) in the Lorenz attractor we were able to map the original Lorenz attractor into something with one hole in the middle — very much like the Rössler attractor, as shown in Fig. 18. This mapping extends from points to orbits, both closed periodic orbits as well as chaotic trajectories, and eventually to branches of the underlying branched manifold. In the general case, suppose we have a chaotic attractor described by a branched manifold with $2n$ branches and an obvious two-fold symmetry. The Lorenz attractor with four branches springs to mind (c.f., Fig. 5). A two-to-one image is constructed by identifying the symmetrically related branches pairwise. The number of branches is halved, as is the number of branch lines and

splitting points. An orbit of minimum period p in the original (cover) attractor is mapped to an orbit of period p in the image system. The image orbit “goes around” either once or twice before closing up. To say this another way, the image orbit has minimum period either p or $p/2$, depending on whether the cover orbit was symmetric or not.

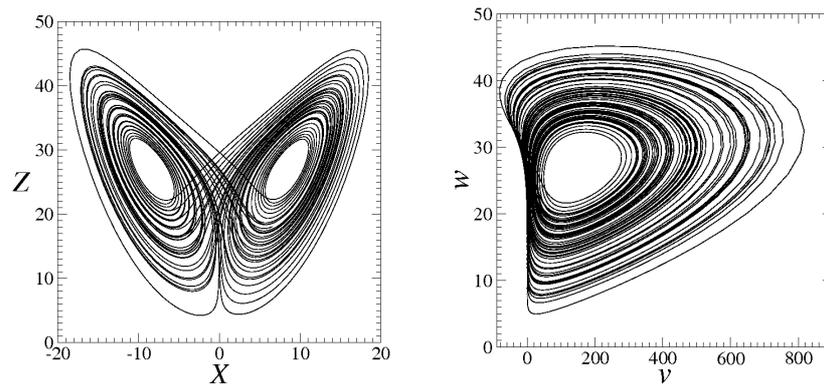


Fig. 18. The Lorenz attractor (left) and its two-to-one image (right). The image is obtained by identifying pairs of points $(X, Y, Z) \leftrightarrow (-X, -Y, +Z)$ in the Lorenz attractor. One convenient way to do this is by introducing a new phase space with coordinates (u, v, w) related to the coordinates (X, Y, Z) through $u = X^2 - Y^2, v = 2XY, w = Z$.

Going in the other direction was yet more exciting. Starting with an image attractor, suppose we wanted to create a two-fold cover with rotational symmetry. Then we have to “lift” the image. The lift can be carried out in many inequivalent ways, depending where we place the rotation axis.⁷⁴ For example, suppose we wanted to create a two-fold cover of the Rössler attractor (2 branches) with rotational symmetry. We could put the rotational axis through the hole in the middle. This gives us a cover with a hole in the middle and four branches organized as a double fold. We could put the rotation axis outside the Rössler attractor. The lift would create two disconnected attractors, each with a hole in the middle and two branches, each of the two identical to the original attractor. We could put the rotation axis between branches 0 and 1 in the Rössler attractor: One way to do this results in a four branch lift that is topologically similar to the Lorenz attractor. Another way of placing the rotation axis between branches 0 and 1 results in an attractor quite different from a Lorenz-like attractor.⁷⁶

All these lifts have four branches. They are structurally stable under perturbation in the location/orientation of the rotation axis. However, if the rotation axis is located in such a way that it intersects the attractor, structurally unstable lifts with six-branches result.⁷⁶

We have used one of these structurally stable lifts to relate sunspot number data N (all positive), which has an approximately 11 year cyclic variation, to the under-

lying magnetic field B , which exhibits an approximately 22 year variation through both positive and negative values.⁷⁵ The usual attempts to relate the two involve making a sign change (by hand) at sunspot minima. We made a planar projection of the data onto the (N, \dot{N}) plane, constructed a lift with two-fold rotational symmetry, and identified one of the cover variables as a surrogate for the underlying magnetic field. The zero-crossings of the magnetic field appeared in a natural way.⁷⁵ The results are shown in Fig. 19. The transformations of the original sunspot data to the time variations of the underlying magnetic field are shown in Fig. 20.

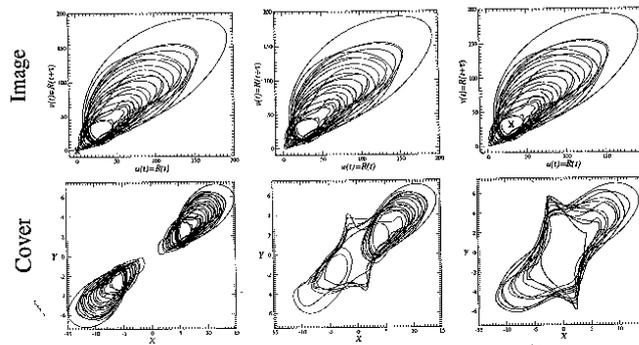


Fig. 19. Top row: Smoothed sunspot data $N(t)$, plotted in the (N, \dot{N}) plane. Bottom row: double covers created with different rotation axis. When the axis is outside the attractor (left column) the cover consists of two separate pieces that do not interact. When the axis intersects with the attractor (center) the cover consists of a single attractor, but regions of positive values are not always succeeded by regions of negative values. When the axis is inside the attractor the double cover has a hole in the middle, but there is a deterministic oscillation between positive values and negative values.⁷⁵

Lifts using larger symmetry groups than the two-fold rotation group can be constructed. The description of many of the exciting things that can happen is a long story that is laid out in a work that brought great joy to us.⁷⁶

Many of the lifts that we considered leave one or more points fixed (rotations around an axis, inversions in a point, reflections in a plane). These point group symmetries act easily on branched manifolds. If $i = 1, 2, 3, \dots$ label the branches in the image, the branches in the cover are labeled $i = 1\alpha, 2\beta, 3\gamma, \dots$, where $\alpha, \beta, \gamma, \dots$ are the operations in the group. A structurally stable lift of an image with b branches by a point group with g group operations has $b \times g$ branches.

Covers of an image attractor with topological entropy h_T have the same topological entropy. The argument is easy — we provide it for a two-fold cover with

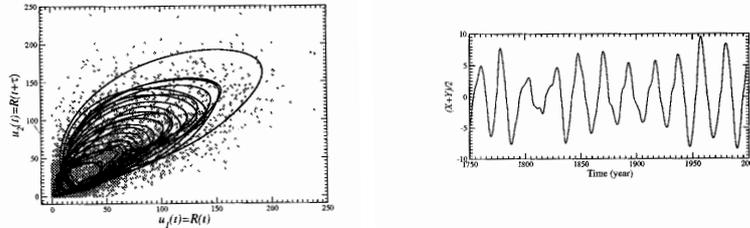


Fig. 20. (Left) Original sunspot data and processed time series, showing an attractor with a hole in the middle. (Right) Time series resulting from a projection of a trajectory in the double cover (c.f. Fig. 19, right column) onto an axis that serves as a surrogate for the underlying magnetic field. This result indicates the strength and polarity of the magnetic field underlying sunspot number variability, one with a 22 year cycle, the other (image) with an 11 year cycle.⁷⁵

rotational symmetry. The lift of a period- p orbit in the image is a symmetric orbit of period $2p$ or two distinct orbits of period p related to each other by symmetry. Assume that the number of orbits of period p grows exponentially in both the cover ($N_C(p) \simeq e^{\gamma p}$) and the image ($N_I(p) \simeq e^{h_T p}$). The number of odd period orbits in the cover is

$$N_C(2p + 1) \simeq e^{\gamma(2p+1)} = 2 \times e^{h_T(2p+1)} \simeq 2N_I(2p + 1) \tag{6}$$

It is a simple exercise to conclude that $\gamma = h_T$ in the limit $2p+1 \rightarrow \infty$. An argument using orbits of both even and odd period in the cover is slightly less simple but leads to the same place: $h_{Cover} = h_{Image} = h_T$. This argument extends without difficulty to covers created using other finite groups.

In addition to symmetries involving point groups, there are symmetries analogous to nonsymmorphic space groups in solid state physics. The simplest such group is the symmetry present when a nonlinear two-dimensional oscillator with inversion symmetry ($\dot{x}_i = f_i(x, y), f_i(-x, -y) = -f_i(x, y)$) is periodically driven: $f_i(x, y) \rightarrow f_i(x, y) + a_i \cos(\omega t + \phi_i)$. In this case the dynamics is invariant under the transformation $(x, y, t) \rightarrow (-x, -y, t + \frac{1}{2}T)$, where $T = 2\pi/\omega$ is the period of the drive. The periodically driven Duffing and van der Pol oscillators share this symmetry.⁴ In these cases it is possible to mod out the order-two symmetry and construct a smaller chaotic attractor that is simpler to analyze.^{4,42} This is done by viewing the attractor from a rotating frame of reference. This frame is called the “van der Pol” frame. In fact, there are two such frames: they counterrotate with angular frequencies $\pm \frac{1}{2}\omega$. There is then a natural question: which to use? Since we are physicists we choose the frame in which the kinetic energy is minimum. The

concepts of average kinetic energy, average rotational energy, and their variation in a series of rotating frames that satisfy simple boundary conditions, is well-defined and natural. These frames are defined by an integer.^{4,78} The results can be seen in Fig. 21.

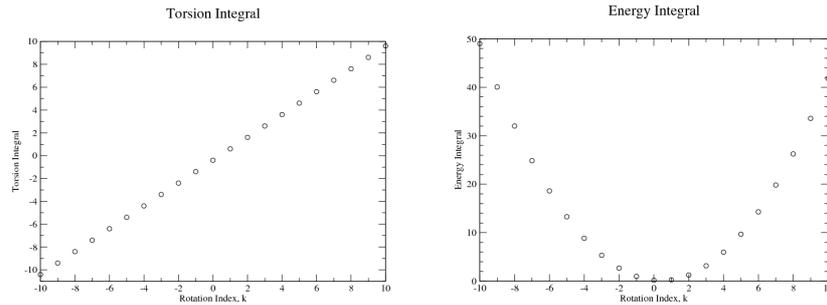


Fig. 21. It is possible to define an average kinetic energy and an average orbital angular momentum for a dynamical system whose phase space is a torus.⁷⁸ These real numbers depend on the representation of the dynamical system: its global torsion n . As the global torsion increases in magnitude, so also do these classical averages. The preferred frame for physicists is the minimum energy frame.

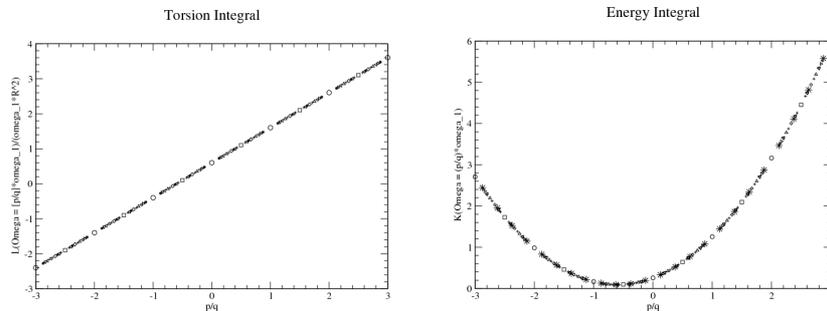


Fig. 22. If a lift is given a fraction (n/p) of a twist per period, then p of these “units” fitted in sequence provide a p -fold cover with global torsion n . The classical integrals depend “smoothly” on the rational fraction n/p .

In the reverse direction, it is possible to lift the image attractor to double covers and p -fold covers, in a very large variety of ways. The multiplicity of possible lifts are distinguished by two integers. The mean energy and mean orbital angular momentum depend on the ratio of these integers, as shown in Fig. 22. Their regularity properties were described in a paper originally entitled “Quantum numbers for strange attractors.” But the referee(s) objected to this title, so it was reluctantly changed to something more prosaic.⁷⁸ An animated lift of the periodically driven van der Pol attractor with three-fold rotational symmetry was constructed

by Timothy Jones and can be enjoyed (“eye-candy”) at his web site.⁷⁹

We were delighted with the results of our cover-image studies. It *was* possible to answer our original question (about Cartan’s theorem) in a way that spoke directly to the close relation between group theory and symmetric dynamical systems. This connection is shown in Fig. 23.

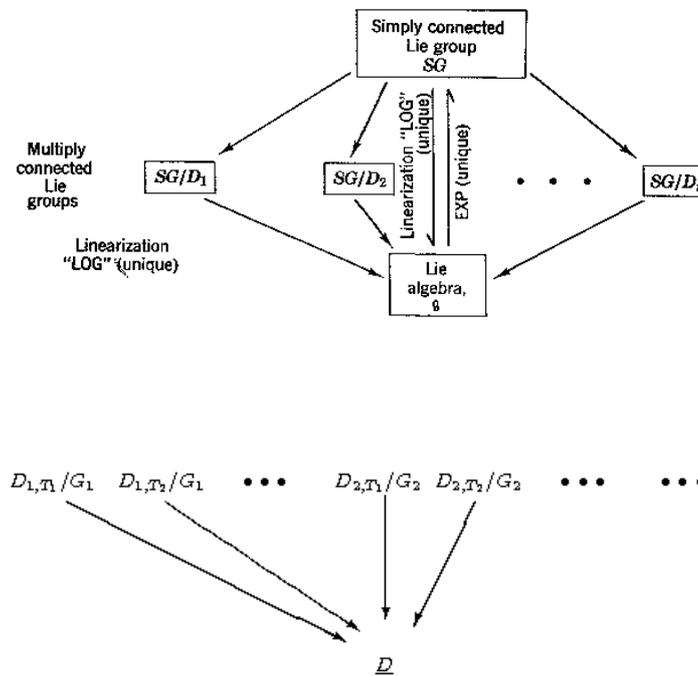


Fig. 23. (top) For Lie groups there is a 1:1 correspondence between Lie algebras and simply connected Lie groups. Each such Lie group is a covering group for *all* Lie groups with the same Lie algebra. These other groups are obtained by “modding out the symmetry”. (bottom) For dynamical systems the relation goes the other way. For each universal image there are many covering dynamical systems. These are distinguished from each other by: the symmetry group G and; the topological index T of the lift.⁷⁶

13. Representation Theory

The first step that needs to be taken for the successful analysis of data taken from a dynamical system behaving chaotically is that a suitable visualization of the system

should be created. If the dynamical system is a set of three ordinary differential equations (viz.: Lorenz or Rössler equations) this is a no-brainer. One only has to watch how the trajectory evolves in the natural three dimensional phase space of coordinates: $(x(t), y(t), z(t)) \in R^3$. If the data are generated in the course of an experiment the situation becomes more exciting. Often the data consists of a single scalar time series $x_i = x(t_i)$. One then hopes to construct a D -dimensional phase space and visualize the trajectory in that phase space. One must construct an *embedding* of the data. Fortunately, this is always possible (in principal) if the original system is finite dimensional ($d < \infty$). This is due to a theorem by Takens⁸⁰ that was exploited by Packard, Crutchfield, Farmer and Shaw⁸¹ and that goes back to Whitney.⁸² The theorem states that if the data are generated by a dynamical system of dimension d an embedding can always be found (is “generic”) in R^D for $D \geq 2d + 1$.

Many embeddings of a scalar time series are possible. The first choice of many is the time delay embedding.^{80,81} This is useful because the signal-to-noise ratio in each component of the embedding is the same. It is less useful because it is not always easy to determine the signs of crossings in two-dimensional projections of time delay embeddings into R^3 . My preferred embedding is the differential embedding. In this case $x \rightarrow \mathbf{y} = (y_1, y_2, y_3)$, with $y_1 = x, y_2 = \dot{x}, y_3 = \ddot{x}$. It is useful because the signs of crossings in projections to the (y_1, y_2) plane are very simple to determine.⁴² It is not useful because the S/N ratio decreases by an order of magnitude (or more) for each higher component of the embedding.

Many inequivalent embeddings of an experimental scalar time series are possible. We emphasize this point by introducing an interesting class of embeddings for the simplest type of dynamical system: a chaotic attractor with a hole in the middle (genus-one strange attractor). Assume we have the scalar time series $x(t)$ and from it we construct the projection into R^2 : $(x(t), y(t) = \dot{x}(t))$. This cannot be an embedding but it could have a hole in the middle. Assume it not only has a hole in the middle, but that a straight line segment attached to a fixed point somewhere inside the hole has the property that the projected flow $(x(t), y(t))$ always strikes the segment from the same side as it is rotated through 2π radians around the fixed point inside the hole. It is then possible to reparameterize the projected trajectory in terms of a rotation angle θ in place of time t : $(x(t), y(t)) \rightarrow (x(\theta), y(\theta))$.

Now introduce the three coordinates (ξ, η, ζ) for a *harmonic knot*:⁸³

$$\mathbf{X}(\theta) = (\xi(\theta), \eta(\theta), \zeta(\theta)) \quad \text{where} \quad \begin{aligned} \xi(\theta) &= \sum_{j=1} A_j \sin(j\theta + \phi_j) \\ \eta(\theta) &= \sum_{j=1} B_j \sin(j\theta + \psi_j) \\ \zeta(\theta) &= \sum_{j=1} C_j \sin(j\theta + \chi_j) \end{aligned} \quad (7)$$

Harmonic knots have periodicity 2π under the parameterization given. Any knot can be given a harmonic parameterization.

Introduce a *repere mobile* for the harmonic knot by constructing the unit tangent, normal, and binormal vectors $(\mathbf{t}(\theta), \mathbf{n}(\theta), \mathbf{b}(\theta))$.⁸⁴ Under suitable not very

restrictive assumptions the data from a chaotic attractor can be embedding using any harmonic knot as a “carrier knot” by

$$x(t) \rightarrow x(\theta) \rightarrow \mathbf{X}(\theta) + x(\theta)\mathbf{n}(\theta) + x'(\theta)\mathbf{b}(\theta) \quad (8)$$

with $' = d/d\theta$. For each knot type there is an embedding of data generated by genus-one dynamics. The embeddings are topologically inequivalent (nonisotopic) because inequivalent knots, the cores of the tori that “carry” the embedding, are not isotopic.

We now come back to the jackpot question that was raised when the topological analysis program was being formulated:³⁹

When you analyze embedded data, what do you learn about the embedding and what do you learn about the dynamical system?

This question struck a chord with me. In Quantum Physics, groups act through their representations. One group can have many inequivalent representations (equivalence is with respect to a change of basis, or similarity transformation). The corresponding question would be: How much can you learn about a group from some/all of its representations.

It seemed reasonable to think that a representation theory of dynamical systems (or their strange attractors) ought to exist⁴ which is spiritually similar to the representation theory of groups/algebras that has found so much use in physics. I discussed this problem with Daniel J. Cross, my graduate student at the time. And presto! After two years of hard work we had a representation theory for strange attractors - at least their low-dimensional varieties.⁸⁵⁻⁸⁸

This theory starts from the natural question: What are the labels for inequivalent representations of a (low-dimensional) strange attractor. By representation labels we mean labels that identify distinct, nonisotopic embeddings of a low-dimensional strange attractor into R^3 . This isn't exactly an easy question (most questions, asked for the first time, aren't easy). So we began with genus-one attractors as a warm-up exercise. In this case we were able to show that there are only three indices. One is parity. The mirror image of a strange attractor in R^3 is diffeomorphic with the original but not isotopic to the original. A second index is knot type. This has been described above in the context of harmonic knots. The third index is global torsion, known very early on from the initial work with Solari.³³ These three indices serve to distinguish all topologically inequivalent (nonisotopic) embeddings of a genus-one dynamics into R^3 .⁷⁰

Now we ask the question: If we raise the embedding dimension by one, mapping $R^3 \rightarrow R^4$, do some of the formerly nonisotopic embeddings into R^3 become equivalent? After all, the representation labels are in some natural sense *obstructions* to isotopy in R^3 . The answer is: Yes, embeddings with different knot types in R^3 all become equivalent in R^4 . Embeddings that are reflected images of each other also

become isotopic in R^4 . Only the global torsion remains an obstruction to isotopy — barely. Embeddings that differ only by a global torsion of 2 become isotopic in R^4 . This means that there are only two inequivalent embeddings into R^4 : Those with even global torsion and those with odd global torsion. In R^5 every embedding is equivalent. The following table^{85,88} shows the progressive extinction of obstructions to isotopy:

Index	R^3	R^4	R^5
Parity	Z_2	—	—
Knot Type	K	—	—
Global Torsion	n	Z_2	—

(9)

In this table K is an index describing knot type (we still haven't a complete handle on this index), Z_2 for parity is ± 1 , and Z_2 for global torsion is $(0, 1)$ or $n \pmod 2$.

A similar result holds for embeddings of genus- g dynamical systems into R^3 . Three indices are required to distinguish among inequivalent embeddings.⁷¹ Raise the embedding dimension by one and only one index remains to describe the residual obstructions to isotopy. In R^5 all embeddings become equivalent.^{85,88}

To answer the jackpot question posed above: Anything learned from analyzing the embedding of a “low-dimensional” dynamical system into R^5 is uniquely about the dynamical system, since all embeddings are equivalent in this space.

14. Pointers to the Future

Topological tools have greatly expanded our understanding of low (=3) -dimensional strange attractors. These attractors are on the borderline of our visual comprehension. They live in three-space but are visualized in 2-space: On the screens of computers, for example. A powerful theorem by Birman and Williams allows us to project attractors down to mostly two-dimensional subspaces. Perhaps more can be said about three-dimensional attractors, but I think the most important things have now been said (I hope I'm wrong!). The one remaining piece of information that would be useful has one foot in R^3 and the other in R^5 . In R^5 there must be some topological signature that exists that can distinguish inequivalent attractors (as opposed to inequivalent embeddings of a single attractor). What is this index and how does it work? This index identifies *mechanism* (c.f. Sec. 9 and Figs. 12 and 13).

A piece of it identifies genus; another piece identifies components of the global Poincaré surface of section; yet another identifies the transition matrix that describes the flow; and yet another describes the stretching-squeezing process that builds up the strange attractor by infinite repetition of the S & S processes. This topological index is the final piece we need for a complete accounting of three-dimensional strange attractors.

Beyond three there is four, and then five, \dots . All my attempts in the last 15 years to extend the topological analysis program to higher dimensions have foundered on one detail. The detail is that Gauss apparently did not extend his knotting thoughts about closed loops for more than a millisecond to more than three dimensions. The Gauss linking number is at the heart of the topological analysis method. It does not extend beyond R^3 . Our analysis methodology therefore does not extend beyond R^3 . We are stuck at the starting line!

It almost seems to me that we may be trying too hard for too much. In low dimensions our goal has evolved to one of determining mechanism. Linking numbers have been a tool to this end. Perhaps we should concentrate more on ends and less on means. This means learn how to classify the stretching and squeezing mechanisms that can operator in R^D ($D > 3$) rather than identifying all the orbits in the attractor, and using them to determine mechanism.

The slides presented at the Birthday Party can be found at.⁸⁹

15. Acknowledgements

I would like to extend heartfelt thanks to the authors of this celebration for their extreme efforts over a one year period to create this conference in the utmost secrecy. This was quite a surprise party.

References

1. M. J. Feigenbaum, Quantitative universality for a class of nonlinear transformations, *J. Stat. Phys.* **19**, 25-52 (1978).
2. M. J. Feigenbaum, The universal metric properties of nonlinear transformations, *J. Stat. Phys.* **21**, 669-706 (1979).
3. M. J. Feigenbaum, Universal behavior in nonlinear systems, *Los Alamos Science* **1**, 4-27 (1980).
4. R. Gilmore and M. Lefranc, *The Topology of Chaos*, NY: Wiley, 2002,
5. Leo Tolstoy, *Anna Karenina*, (tr. Constance Garnett), NY: Barnes and Noble, 2000.
6. James Gleick, *Chaos, The Making of a New Science*, Viking, Penguin, 1987.
7. M. Giglio, S. Musazzi, and U. Perini, Transition to chaotic behavior via a reproducible sequence of period-doubling bifurcations, *Phys. Rev. Lett.* **47**, 243-246 (1981).
8. A. Libchaber, C. Laroche, and S. Fauve, Period doubling cascade in Mercury, a quantitative measurement, *Le Journal de Physique - Letters* **43**, L211-L216 (1982).
9. R. H. Simoyi, A. Wolf, and H. L. Swinney, One-dimensional dynamics in a multicomponent chemical reaction, *Phys. Rev. Lett.* **49**, 245-248 (1982).
10. J. C. Roux, R. H. Simoyi, and H. L. Swinney, Observation of a strange attractor, *Physica* **D8**, 257-266 (1983).
11. K. Coffman, W. D. McCormick, Z. Noszticzius, R. H. Simoyi, and H. L. Swinney, Universality, multiplicity and the effect of iron impurities in the Belousov-Zhabotinskii reaction, *J. Chem. Phys.* **86**, 119-129 (1987).
12. J. Testa, J. Perez, and C. Jeffries, Evidence for universal chaotic behavior of a driven nonlinear oscillator, *Phys. Rev. Lett.* **48**, 1217-1220 (1982).
13. F. T. Arecchi, R. M. Meucci, G. Puccioni, and J. R. Tredicce, Experimental evidence

- of subharmonic bifurcations, multistability, and turbulence in a Q-switched gas laser, *Phys. Rev. Lett.* **49**, 245-248 (1982).
14. R. S. Gioggia and N. B. Abraham, Routes to chaotic output from a single mode, dc-excited laser, *Phys. Rev. Lett.* **51**, 650-653 (1983).
 15. G. P. Puccioni, A. Poggi, W. Gadomski, J. R. Tredicce, and F. T. Arecchi, Measurement of the formation and evolution of a strange attractor in a laser, *Phys. Rev. Lett.* **55**, 339-342 (1985).
 16. J. R. Tredicce, F. T. Arecchi, G. P. Puccioni, A. Poggi, and W. Gadomski, Dynamic behavior and onset of low-dimensional chaos in a modulated homogeneously broadened single-mode laser: Experiments and theory, *Phys. Rev.* **A34**(3), 2073-2081 (1986).
 17. P. Grassberger, On the Hausdorff dimension of fractal attractors, *J. Stat. Phys.* **26**(1), 173-179 (1981).
 18. P. Grassberger, Generalized dimensions of strange attractors, *Phys. Lett.* **97**(6), 227-230 (1983).
 19. P. Grassberger and I. Procaccia, Measuring the strangeness of strange attractors, *Physica* **D9**, 189-208 (1983).
 20. P. Grassberger and I. Procaccia, Characterization of strange attractors, *Phys. Rev. Lett.* **50**(5), 346-349 (1983).
 21. J.-P. Eckmann and D. Ruelle, Ergodic theory of chaos and strange attractors, *Rev. Mod. Phys.* **57**, 617-656 (1985).
 22. A. Wolf, J. B. Swift, H. L. Swinney, and J. A. Vastano, Determining Lyapunov exponents from a time series, *Physica* **D16**, 285-317 (1985).
 23. S. Sato, M. Sano, and Y. Sawada, Practical methods of measuring the generalized dimension and largest Lyapunov exponent in high dimensional chaotic systems, *Prog. Theor. Phys.* **77**(1), 1-5 (1987).
 24. R. Stoop and J. Parisi, Calculation of Lyapunov exponents avoiding spurious elements, *Physica* **D50**, 89 (1991).
 25. M. T. Rosenstein, J. J. Collins, and C. J. De Luca, A practical method for calculating largest Lyapunov exponents from small data sets, *Physica* **D65**, 117-134 (1993).
 26. D. Ruelle, *Proc. Roy. Soc. London, Ser. A* **427**, 241 (1990).
 27. E. N. Lorenz, Deterministic nonperiodic flow, *J. Atmos. Sci.* **20**, 130-147 (1963).
 28. O. E. RöSSLer, An equation for continuous chaos, *Phys. Lett.* **A57**(5), 397-398 (1976).
 29. H. Poincaré, *Les Methodes nouvelle de la mécanique céleste*, Paris: Gauthier-Villars, 1892.
 30. D. Auerbach, P. Cvitanović, J.-P. Eckmann, G. Gunaratne, and I. Procaccia, Exploring chaotic motion through periodic orbits, *Phys. Rev. Lett.* **58**, 2387-2390 (1987).
 31. D. J. Lathrop and E. J. Kostelich, Characterization of an experimental strange attractor by periodic orbits, *Phys. Rev.* **A40**, 4028-4031 (1989).
 32. J.-P. Eckmann, S. O. Kamphorst, and D. Ruelle, Recurrence plots of dynamical systems, *Europhys. Lett.* **5**, 973-977 (1987).
 33. H. G. Solari and R. Gilmore, Relative rotation rates for driven dynamical systems, *Phys. Rev.* **A37**, 3096-3109 (1988).
 34. N. B. Tuffillaro, H. G. Solari, and R. Gilmore, Relative rotation rates: Fingerprints for strange attractors, *Phys. Rev.* **A41**, 5717-5720 (1990).
 35. G. B. Mindlin, X.-J. Hou, H. G. Solari, R. Gilmore, and N. B. Tuffillaro, Classification of strange attractors by integers, *Phys. Rev. Lett.* **64**, 2350-2353 (1990).
 36. J. S. Birman and R. F. Williams, Knotted periodic orbits in dynamical systems I: Lorenz's equations, *Topology* **22**(1), 47-82 (1983).
 37. J. S. Birman and R. F. Williams, Knotted periodic orbits in dynamical systems II: Knot holders for fibered knots, *Contemporary Mathematics* **20**, 1-60 (1983).

38. R. F. Williams, personal communication to the author.
39. G. B. Mindlin, H. G. Solari, M. A. Natiello, R. Gilmore, and X.-J. Hou, Topological analysis of chaotic time series data from the Belousov-Zhabotinskii reaction, *J. Nonlinear Science* **1**, 147-173 (1991).
40. J. Milnor and W. Thurston, *On iterated maps of the interval Dynamical systems, College Park, MD, 1986-87 Lecture Notes in Math.*, **1342**, Berlin: Springer-Verlag, 1988: pp. 465-563.
41. M. Lefranc, D. Hennequin, and P. Glorieux, Homoclinic chaos in a laser containing a saturable absorber, *Phys. Lett. A* **163**, 239-249 (1992).
42. R. Gilmore, topological analysis of chaotic dynamical systems, *Reviews of Modern Physics* **70**, 1455-1530 (1998).
43. A. Amon and M. Lefranc, Topological signature of deterministic chaos in short nonstationary signals from an optical parametric oscillator, *Phys. Rev. Lett.* **92**(9), 094101 (2004).
44. D. Hennequin, F. de Tomasi, B. Zambon, and E. Arimondo, Homoclinic orbits and cycles in the instabilities of a laser with saturable absorber, *Phys. Rev.* **A37**, 2243-2246 (1988).
45. F. A. Papoff, A. Fioretti, E. Arimondo, G. B. Mindlin, H. G. Solari, and R. Gilmore, Structure of chaos in the laser with saturable absorber, *Phys. Rev. Lett.* **68**, 1128-1131 (1992).
46. T. Shimizu and N. Morioka, On the bifurcation of a symmetric limit cycle to an asymmetric one in a simple model, *Physics Letters* **A76**, 201-204 (1980).
47. N. B. Tufillaro, T. A. Abbott, and J. P. Reilly, *An Experimental Approach to Nonlinear Dynamics and Chaos*, Reading, MA: Addison-Wesley, 1992.
48. G. B. Mindlin, R. Lopez-Ruiz, H. G. Solari, and R. Gilmore, Horseshoe implications, *Phys. Rev.* **E48**, 4297-4304 (1993).
49. T. Hall, Weak Universality in two-dimensional transitions to chaos, *Phys. Rev. Lett.* **71**, 58-61 (1992).
50. T. Hall and A. de Carvalho, How to prune a horseshoe, *Nonlinearity* **15** R19-R68 (2002).
51. T. Hall and A. de Carvalho, The forcing relation for horseshoe braid types, *Experimental Math.* **11** 271-288 (2002).
52. M.A. Natiello and H.G. Solari, Remarks on braid theory and the characterisation of periodic orbits, *Journal of Knot Theory and its Ramifications* **3**(4), 511-529, (1994).
53. M. Lefranc, Alternative determinism principle for topological analysis of chaos, *Phys. Rev.* **E74**, 035202(R) (2006).
54. P. Collins, Forcing relations for homoclinic orbits of the Smale horseshoe map, *Experimental Mathematics* **14**(1), 75-86 (2005).
55. J. L. W. McCallum and R. Gilmore, A geometric model of the driven Duffing oscillator, *Int. J. Bifurcation and Chaos* **3**, 685-691 (1993).
56. R. Gilmore and J. W. L. McCallum, Superstructure in the bifurcation diagram of the Duffing oscillator, *Phys. Rev.* **E51**, 935-956 (1995).
57. R. Gilmore, C. Letellier, and M. Lefranc, Chaos - Topology, *Scholarpedia* **3**(7):4592 (2007).
58. G. Boulant, M. Lefranc, S. Bielawski, and D. Derozier, Horseshoe templates with global torsion in a driven laser, *Phys. Rev.* **E55**(5), 5082-5091 (1997).
59. R. Gilmore, X. Pei, and F. Moss, Topological analysis of chaos in neural spike train bursts, *Chaos* **9**(3), 812-817 (1999).
60. R. Gilmore and X. Pei, The topology and organization of unstable periodic orbits in Hodgkin-Huxley models of receptors with subthreshold oscillations, *Handbook of Bi-*

- ological Physics, Vol. 4, Neuro-informatics, Neural Modeling*, (F. Moss and S. Gielen, Eds.), Amsterdam: North Holland, 2001, pp. 155-203.
61. J. Used and J. C. Martin, Reverse horseshoe and spiral templates in an Erbium doped fiber laser. *Phys. Rev.* **E79**, 046213 (2009).
 62. J. Used and J. C. Martin, Multiple topological structures of chaotic attractors ruling the emission of a driven laser, *Phys. Rev.* **E82**, 016218 (2010).
 63. G. Byrne, R. Gilmore, and C. Letellier, Distinguishing between folding and tearing mechanisms in strange attractors, *Phys. Rev.* **E70**, 056214 (2004).
 64. C. Letellier, T. D. Tsankov, G. Byrne, and R. Gilmore, Large scale structural reorganization of strange attractors, *Phys. Rev.* **E72**, 026212 (2005).
 65. T. D. Tsankov and R. Gilmore, Strange attractors are classified by bounding tori, *Phys. Rev. Lett.* **91**(13), 134104 (2003).
 66. T. D. Tsankov and R. Gilmore, Topological aspects of the structure of chaotic attractors in \mathbb{R}^3 , *Phys. Rev.* **E69**, 056206 (2004).
 67. J. Katriel and R. Gilmore, Entropy of bounding tori, *Entropy Online* **12**, 066220 (2010).
 68. V. I. Arnold, *Ordinary Differential Equations*, Cambridge, MA: MIT Press, 1980.
 69. M. A. Aziz-Alaoui, Differential equations with multispiral attractors, *Int. J. Bifurcation Chaos Appl. Sci. Eng.* **9**, 1009-1039 (1999).
 70. N. Romanazzi, M. Lefranc, and R. Gilmore, Embeddings of low-dimensional strange attractors: Topological invariants and degrees of freedom, *Physical Review* **E75**, 066214 (2007).
 71. R. Gilmore, C. Letellier, and N. Romanazzi, Global topology from an embedding, *J. Phys. A: Math. Theor.* **40**, 13,291 - 13,297 (2007) .
 72. R. Gilmore, *Lie Groups, Lie Algebras, and Some of Their Applications*, NY: Wiley, 1974.
 73. C. Letellier and R. Gilmore, Covering dynamical systems: Two-fold covers, *Phys. Rev.* **E63**, 016206 (2000).
 74. C. Letellier and R. Gilmore, Dressed symbolic dynamics, *Phys. Rev.* **E67**, 036205 (2003).
 75. C. Letellier, L. A. Aguirre, J. Maquet, and R. Gilmore, Evidence for low-dimensional chaos in sunspot cycles, *Astronomy and Astrophysics* **449**, 379-387 (2006).
 76. R. Gilmore and C. Letellier, *The Symmetry of Chaos*, Oxford: University Press, 2008.
 77. C. Letellier and R. Gilmore, Symmetry groups for 3D dynamical systems, *Journal of Physics* **A40**, 5597-5620 (2007).
 78. R. Gilmore, Two-parameter families of strange attractors, *Chaos* **17**, 013104 (2007).
 79. Tim Jones: <http://lagrange.physics.drexel.edu/flash/vanderpol/>
 80. F. Takens, Detecting strange attractors in turbulence, in: *Lecture Notes in Mathematics*, Vol. 898, NY: Springer, 1981.
 81. N. H. Packard, J. P. Crutchfield, J. D. Farmer, and R. S. Shaw, Geometry from a time series, *Phys. Rev. Lett.* **45**, 712-716 (1980).
 82. H. Whitney, Differentiable manifolds, *Annals of Mathematics* **37**(3), 645-680 (1936).
 83. <http://www.math.uiowa.edu/~jsimon/aaron.html>
 84. D. J. Struik, *Lectures on Classical Differential Geometry*, Reading, MA: Addison-Wesley, 1950.
 85. Daniel J. Cross and R. Gilmore, Representation theory for strange attractors, *Phys. Rev.* **E80**(1), 056207 (2009).
 86. Daniel J. Cross and R. Gilmore, Differential embedding of the Lorenz attractor, *Physical Review* **E81**, 066220 (2010).
 87. Daniel J. Cross and R. Gilmore, Equivariant differential embeddings, *Journal of Math-*

- ematical Physics **51**, 092706 (2010).
88. Daniel J. Cross and R. Gilmore, Complete set of representations for dissipative chaotic three-dimensional dynamical systems, *Physical Review* **E82**, 056211 (2010).
 89. <http://einstein.drexel.edu/~bob/Presentations/birthday70.pdf>