Volume XXIX 2026
MBNA Publishing House Constanta 2026

Scientific Bulletin of Naval Academy

SBNA PAPER • OPEN ACCESS

Ulam's Method and the Transfer Operator: A Numerical Path to Hausdorff Dimension in Dynamical Systems

To cite this article: Alexandru-Ionel Toma, Scientific Bulletin of Naval Academy, Vol. XXIX 2026, pg. 160-168.

Submitted: 27.04.2026
Revised: 31.07.2026
Accepted: 04.08.2026

Available online at www.anmb.ro

ISSN: 2392-8956; ISSN-L: 1454-864X

doi: 10.21279/1454-864X-26-I1-015

Ulam's Method and the Transfer Operator: A Numerical Path to Hausdorff Dimension in Dynamical Systems

Alexandru-Ionel Toma
Faculty of Automation and Computers, National University of Science and Technology POLITEHNICA, Bucharest, Romania
e-mail: toma.alexandru87@yahoo.com

Abstract. This article presents a unified numerical framework connecting three fundamental concepts in dynamical systems theory: the transfer operator, Ulam's discretization method, and Hausdorff measure. We begin by introducing the transfer operator (Perron–Frobenius operator) as the natural object governing the evolution of probability densities under a chaotic map. Recognizing that this operator acts on an infinite-dimensional space, we describe Ulam's classical finite-box approximation, which reduces the problem to computing the stationary distribution of a finite Markov chain. The central contribution of this exposition is the synthesis of these ideas with the thermodynamic formalism: by considering a parameterized family of transfer operators with geometric weights, we define the pressure function, whose unique zero yields the Hausdorff dimension of the invariant set. Ulam's method applied to this parameterized operator provides a fully numerical pathway to compute fractal dimensions of strange attractors and repellers. We illustrate the theory with explicit examples, discuss numerical implementation challenges, and survey current frontiers including non-uniformly hyperbolic systems and connections to data-driven methods.
Keywords: Transfer operator, Ulam's method, Hausdorff dimension, thermodynamic formalism, invariant measures, numerical approximation.

1 Introduction

How can we assign a numerical size to the intricate sets that often appear in chaotic dynamical systems? Standard Euclidean geometry does not provide the right tools for these objects. Hausdorff measure and dimension were developed precisely for this purpose. However, calculating the Hausdorff dimension for a particular system is still a difficult computational problem.

There are two main parts to this difficulty. The first is statistical. The geometry of an attractor alone does not tell us how often a typical orbit visits different regions of that set. To study this, we shift our attention from single trajectories to evolving probability distributions. This is done using the *transfer operator* (also called the Perron–Frobenius operator). Its invariant functions give the statistical steady state of the system.

The second part is geometric. The geometry and the statistics are not separate. In thermodynamic formalism, they are connected through a family of transfer operators that depend on a parameter. From these operators, one defines a pressure function. The Hausdorff dimension of the invariant set is then given by the unique point where this pressure equals zero. Unfortunately, these operators act on infinite-dimensional spaces, and exact solutions are only available for very simple cases.

This is where numerical methods come in. In 1960, Ulam proposed a simple discretization: divide the phase space into small boxes and replace the infinite-dimensional operator by a finite matrix of transition probabilities. This *Ulam method* is commonly used to approximate invariant densities. More importantly for our purpose, it can also be applied to the parameterized operators. This gives us a way to compute the pressure function numerically and, consequently, to estimate the Hausdorff dimension.

In this article, we describe this numerical approach in three parts. First, we define the transfer operator and explain how it describes the evolution of distributions. Second, we present Ulam's discretization and discuss its basic properties. Third, we introduce the parameterized transfer operator and show how the pressure function leads to a numerical estimate of the Hausdorff dimension. The aim is to make clear how these three ideas work together to provide a practical method for computing fractal dimensions.

2 From Orbits to Distributions: The Transfer Operator

2.1 From Points to Distributions

We consider a discrete-time dynamical system. This consists of a phase space X (often an interval or a more general space) and a map $T : X \to X$. Starting from an initial point x_0, one obtains an orbit by iterating the map:

$$x_0, \quad x_1 = T(x_0), \quad x_2 = T^2(x_0), \quad \dots$$

where T^n means applying T a total of n times.

A standard approach in dynamics is to study the qualitative behavior of such orbits. One may ask whether an orbit converges to a fixed point, becomes periodic, or behaves chaotically. However, this pointwise view has a major limitation in practice. In chaotic systems, initial conditions are never known exactly, and even tiny errors grow quickly. Moreover, a single trajectory, no matter how long, only gives a sparse picture of the long-term behavior. A better approach is to shift attention from individual points to entire distributions. Instead of asking where a specific point goes, one can ask: given an initial probability distribution over the phase space, how does this distribution evolve under the map?

This statistical viewpoint is useful not only because it handles uncertainty, but also because it describes the long-term averages that are observable in experiments. For many chaotic systems, a typical orbit spends different amounts of time in different regions. The limiting distribution, when it exists and is independent of the starting point (except for a negligible set), is called the physical or natural invariant measure. The main task is to find a mathematical way to describe how distributions evolve.

2.2 The Transfer Operator

The evolution of distributions is formalized by the *transfer operator* (sometimes called the Perron–Frobenius operator). Let $\mathcal{M}(X)$ be the space of probability measures on X. The operator

$\mathcal{P}$ is defined by the relation

$$(\mathcal{P}\mu)(A) := \mu\big(T^{-1}(A)\big) \tag{1}$$

for any measurable set $A \subset X$. The interpretation is simple: the mass that ends up in A after one iteration is exactly the mass that was in the preimage $T^{-1}(A)$ before the iteration.

In many situations it is convenient to work with density functions rather than abstract measures. Suppose there is a reference measure m (usually the Lebesgue measure) and that μ has density f with respect to m, i.e., $d\mu = f\,dm$. If the map T is nonsingular (meaning $m(A) = 0$ implies $m(T^{-1}(A)) = 0$), then $\mathcal{P}\mu$ also has a density, which we write as $\mathcal{P}f$. This density is defined by the change of variables formula

$$\int_A (\mathcal{P}f)(x)\,dx = \int_{T^{-1}(A)} f(x)\,dx \tag{2}$$

for all measurable sets A. This equation simply expresses conservation of probability.

For one-dimensional expanding maps with a finite number of branches, the operator can be written in an explicit pointwise form. If $T^{-1}(x) = \{y_1, \ldots, y_k\}$, then

$$(\mathcal{P}f)(x) = \sum_{y \in T^{-1}(x)} \frac{f(y)}{|T'(y)|}. \tag{3}$$

The term $1/|T'(y)|$ accounts for how much the map stretches or compresses space near the preimage y. In higher dimensions, the derivative is replaced by the Jacobian determinant.

The fixed points of the transfer operator give the invariant measures. A density f_* is called invariant if

$$\mathcal{P}f_* = f_*. \tag{4}$$

For expanding or hyperbolic systems, there is a particular invariant density that is physically relevant, called the SRB measure (after Sinai, Ruelle, and Bowen). It has the property that for typical starting points x,

$$\lim_{n\to\infty} \frac{1}{n} \sum_{k=0}^{n-1} \phi(T^k x) = \int_X \phi(x) f_*(x)\,dx \tag{5}$$

for any continuous observable ϕ. In other words, the invariant density describes how a typical orbit is distributed in the long run.

2.3 A Concrete Example: The Doubling Map

To see how the transfer operator works in a simple case, consider the doubling map on the unit interval:

$$T(x) = 2x \pmod 1, \quad x \in [0, 1).$$

This map has two branches, each with derivative 2. For a constant density $f(x) = 1$, we have two preimages for each point x, namely $y_1 = x/2$ and $y_2 = (x+1)/2$. The formula gives

$$(\mathcal{P}1)(x) = \frac{1}{2} + \frac{1}{2} = 1.$$

Thus the constant density is invariant. This confirms that the uniform distribution is the natural invariant measure for the doubling map. A typical orbit is uniformly distributed over the interval.

For more complicated, non-linear maps, the invariant density is not constant and cannot usually be found by hand. In such cases, one has to resort to numerical methods, which is the subject of the next part.

3 Ulam's Method: Discretizing the Infinite-Dimensional Operator

3.1 The Ulam Discretization

We saw in the previous section that the invariant density is given by the leading eigenfunction of the transfer operator. For simple examples like the doubling map, this density is just a constant. For most realistic systems, however, the invariant density is not known in closed form. The main issue is that the transfer operator acts on an infinite-dimensional space of functions. To carry out any computation, we must replace this operator by a finite matrix approximation. Ulam proposed a simple way to do this in 1960.

The idea is to partition the phase space into finitely many cells and approximate the dynamics by a finite-state Markov chain. The construction has three steps.

First, we divide the phase space X into n disjoint cells:

$$X = \Delta_1 \cup \Delta_2 \cup \cdots \cup \Delta_n, \quad \Delta_i \cap \Delta_j = \varnothing \text{ for } i \neq j.$$

In one dimension these are intervals; in higher dimensions they are boxes or other simple shapes. For a uniform grid with N subdivisions per dimension, we have $n = N^d$ cells. The cell diameter h determines the accuracy of the approximation, which we expect to improve as h tends to zero.

Second, we build an $n \times n$ transition matrix P whose entries p_{ij} are the probabilities that a point chosen at random in cell Δ_i lands in cell Δ_j after one application of the map. These entries are given by

$$p_{ij} = \frac{m\big(\Delta_i \cap T^{-1}(\Delta_j)\big)}{m(\Delta_i)}, \tag{6}$$

where m is the reference measure (usually Lebesgue measure). The numerator is the size of the portion of Δ_i that maps into Δ_j.

Third, we interpret P as a row-stochastic matrix. Each row sums to one because $\bigcup_j T^{-1}(\Delta_j) = X$ (up to sets of measure zero). This matrix defines a Markov chain on the n cells. The steady-state distribution of this chain is a row vector $\pi = (\pi_1, \ldots, \pi_n)$ satisfying

$$\pi P = \pi, \quad \sum_{i=1}^{n} \pi_i = 1, \quad \pi_i \geq 0. \tag{7}$$

The value π_i gives the approximate fraction of time that a typical orbit spends in cell Δ_i. The corresponding piecewise constant density is

$$\tilde{f}_*(x) = \sum_{i=1}^{n} \frac{\pi_i}{m(\Delta_i)} \mathbf{1}_{\Delta_i}(x), \tag{8}$$

where $\mathbf{1}_{\Delta_i}$ is the indicator function of cell i. The normalization ensures that the total integral of $\tilde{f}_*$ is one.

3.2 Projection Interpretation and Convergence

There is a standard mathematical way to view the construction above. Let $\mathcal{V}_n$ be the space of piecewise constant functions on the partition. Define a projection operator Π_n that averages a function over each cell:

$$(\Pi_n f)(x) = \frac{1}{m(\Delta_i)} \int_{\Delta_i} f(y)\, dy \quad \text{for } x \in \Delta_i.$$

The Ulam matrix P is then the matrix representation of the operator $\mathcal{P}_n = \Pi_n \mathcal{P}\big|_{\mathcal{V}_n}$. In other words, we apply the true transfer operator to a piecewise constant function and then average the result over each cell. This means we give up some fine-scale detail to make the problem computable.

The hope is that as the partition becomes finer (i.e., as n increases), the finite-dimensional approximation improves and approaches the true operator in a suitable sense. It has been shown that for many expanding and uniformly hyperbolic systems, the leading eigenvector of the Ulam matrix converges to the true invariant density. In practice, the method works well for one-dimensional systems and for some higher-dimensional cases, although the convergence can be slow if the invariant density has strong singularities. We will revisit this point in the next part when we extend the method to parameterized operators.

4 Parameterized Operators and the Pressure Function: The Path to Hausdorff Dimension

4.1 Parameterized Operators and Pressure

In the previous sections, we saw how the transfer operator gives the statistical behavior of a system through its invariant density. However, an invariant set such as a strange attractor or a repeller also has geometric structure. These sets are often fractal: they cannot be measured properly using length, area, or volume. The right tool for this is Hausdorff measure, and the main quantity of interest is the Hausdorff dimension. We now explain how the transfer operator can be adapted to compute this dimension.

The connection comes from thermodynamic formalism, developed by Bowen, Ruelle, and Sinai. The idea is to introduce a parameter t into the transfer operator. The geometry of the invariant set is then encoded in the spectral properties of this one-parameter family.

We restrict attention to expanding repellers. These are compact invariant sets Λ on which the map T is expanding: there is a constant $\sigma > 1$ such that $|T'(x)| \geq \sigma$ for all $x \in \Lambda$. The map is also locally invertible, meaning each point has finitely many preimages in Λ. A typical example is the Julia set of a polynomial such as $T(x) = x^2 - 2$ for certain parameter values.

For such systems, we define a family of operators $\mathcal{L}_t$ for $t \in \mathbb{R}$ by

$$(\mathcal{L}_t f)(x) = \sum_{y \in T^{-1}(x)} \frac{f(y)}{|T'(y)|^t}. \tag{9}$$

When $t = 1$, this is just the usual transfer operator. For other values of t, the operator applies a weight that depends on the local expansion rate. The term $|T'(y)|^{-t}$ penalizes preimages where the map expands a lot (large derivative) when $t > 0$, and favors them when $t < 0$. This weighting is directly related to the way Hausdorff measure is defined: when covering a self-similar set, each piece is scaled by a factor that depends on the contraction ratio, and the exponent determines whether the total measure is zero, finite, or infinite.

The pressure function $P(t)$ is defined as the logarithm of the spectral radius of $\mathcal{L}_t$:

$$P(t) := \log\big(\rho(\mathcal{L}_t)\big). \tag{10}$$

For expanding maps, the spectral radius is a dominant eigenvalue. The pressure has a few basic properties: it is strictly decreasing in t, convex, and tends to $+\infty$ as $t \to -\infty$ and to $-\infty$ as $t \to +\infty$. The key result, due to Bowen and Ruelle, is the following:

Theorem 4.1 (Bowen–Ruelle)**.** *Let Λ be a repeller for an expanding C^1 map T. The Hausdorff dimension d_H of Λ is the unique real number t such that*

$$P(t) = 0. \tag{11}$$

Thus, instead of covering the fractal with balls and computing the dimension directly, we can find it by locating the zero of the pressure function. This is a spectral problem: we need to compute the spectral radius of $\mathcal{L}_t$ as a function of t, and then solve for the root.

To see why this works, consider the middle-thirds Cantor set. The associated expanding map has two inverse branches, each with contraction ratio $1/3$, so $|T'| = 3$ on each branch. The operator $\mathcal{L}_t$ acts on constant functions as

$$\mathcal{L}_t(1) = 3^{-t} + 3^{-t} = 2 \cdot 3^{-t}.$$

The pressure is $P(t) = \log(2 \cdot 3^{-t}) = \log 2 - t \log 3$. Setting $P(t) = 0$ gives

$$t = \frac{\log 2}{\log 3},$$

which is exactly the Hausdorff dimension of the Cantor set. The calculation is simple because the contraction is uniform. For non-uniform systems, the pressure equation becomes more complex, and the spectral radius of $\mathcal{L}_t$ gives the appropriate weighted average.

The thermodynamic formalism and the associated pressure functions are not limited to uniformly expanding repellers. They are also central to the metrical theory of continued fraction algorithms, where the transfer operator approach yields precise information about the Hausdorff dimension of exceptional sets; see [8] for a detailed exposition.

4.2 Numerical Computation of Hausdorff Dimension

Since the pressure is the logarithm of the spectral radius of $\mathcal{L}_t$, and since Ulam's method gives a finite approximation to the transfer operator, we can extend the same idea to the parameterized family. This gives a practical numerical algorithm.

The procedure is as follows. Given a partition of X into n cells, we build a matrix $P_n(t)$ whose entries are weighted transition probabilities. For each sample point in a cell, we apply the map, record the destination cell, and accumulate the weight $|T'(y)|^{-t}$ instead of a simple count. The entry $p_{ij}(t)$ is the average of these weights over all samples in cell Δ_i. Once the matrix is constructed, we compute its spectral radius $\rho_n(t)$, typically using power iteration or a similar method. The Ulam approximation to the pressure is then

$$P_n(t) = \log \rho_n(t). \tag{12}$$

The Hausdorff dimension is approximated by the unique t such that $P_n(t) = 0$. Since $P_n(t)$ is decreasing, we can find the root by bisection or Newton's method.

The following pseudocode summarizes the algorithm:

```
function compute_Hausdorff_dimension(T, X, n_cells, t_min, t_max, tol):
    partition X into n cells
    for each t in a bracket [t_min, t_max]:
        build Ulam matrix P_n(t) using sample points weighted by |T'|^{-t}
        compute spectral radius rho = power_iteration(P_n(t))
        pressure = log(rho)
    find root d_H such that pressure(d_H) = 0 via bisection
    return d_H
```

In practice, there are several numerical issues to be aware of. Near the root, the pressure is sensitive to small errors in the matrix entries, so the computation can be ill-conditioned. Monte Carlo sampling introduces statistical noise, especially if the number of sample points per cell is too small. The matrix size grows rapidly with the dimension of the phase space, which limits the method to moderate dimensions unless sparse or adaptive schemes are used. Finally, if the map has critical points where the derivative vanishes, the weights $|T'|^{-t}$ blow up for $t > 0$, requiring special treatment. Despite these difficulties, the method has been applied successfully to a range of systems, including Julia sets and Hénon-type attractors.

5 Conclusions and Future Directions

5.1 Summary and Extensions

Let us briefly summarize what we have covered. In Section 2 we introduced the transfer operator and explained how it describes the evolution of probability densities. The main point was that instead of studying individual trajectories, we can study distributions, and the invariant density of the transfer operator gives the long-term statistical behavior of the system. In Section 3 we presented Ulam's discretization method, which replaces the infinite-dimensional transfer operator by a finite matrix. This makes the problem computable, although the accuracy depends on the fineness of the partition. In Section 4 we extended the transfer operator to a one-parameter family by introducing weights depending on the local expansion rate. The spectral radius of this parameterized operator defines the pressure function, and the Bowen–Ruelle theorem tells us that the Hausdorff dimension of the invariant set is the unique zero of this pressure. Applying Ulam's discretization to the parameterized operator gives a practical numerical method for computing fractal dimensions.

There are several ways in which this framework can be extended. The classical theory assumes uniform expansion, but many systems of interest are not uniformly hyperbolic. Examples include intermittent maps with indifferent fixed points and the Hénon map. One approach to handle such systems is to use induced maps: one constructs a first-return map to a suitable subset where the dynamics become uniformly expanding, applies the Ulam method to this induced system, and then translates the results back. The pressure equation becomes more involved, but the basic idea remains the same.

Another direction concerns multifractal spectra. The leading eigenvalue of the parameterized operator gives the Hausdorff dimension of the whole invariant set. However, fractal sets are often not homogeneous: different regions may have different local scaling exponents. The multifractal spectrum describes the dimensions of the subsets where the local dimension takes a given value. This spectrum can be obtained from the subdominant eigenvalues of the operator via a Legendre transform. Ulam's method can be extended to approximate these eigenvalues as well.

5.2 Practical Considerations and Future Work

For those who wish to implement the numerical method described here, a few practical points are worth mentioning. It is advisable to start with one-dimensional systems where the results can be checked against known formulas. Uniform grids are simple to implement, but adaptive partitions can improve accuracy significantly for systems with non-uniform contraction rates. The Ulam matrix is often sparse, so storing it in a sparse format and using iterative eigenvalue solvers can save memory and time. Before applying the method to a new system, it is helpful to test it on cases with known Hausdorff dimensions, such as the middle-thirds Cantor set, to calibrate the sampling density and convergence criteria. If the map has singularities or discontinuities in the

derivative, one should use finer sampling in those regions or use higher-order quadrature instead of simple Monte Carlo.

Despite the progress made, several open problems remain. A particularly natural next step is to apply the Ulam-pressure framework to the wide class of continued fraction expansion algorithms. For the Gauss map and its generalizations, the pressure function is typically studied via an induced map approach, where the indifferent fixed point is dealt with by considering the first-return map to a suitable subset. Developing efficient numerical schemes for these non-uniformly expanding systems, and validating them against known analytic results, is a promising direction for future research. The recent monograph by Sebe and Lascu [8] provides a comprehensive background on the metrical and ergodic theory of such algorithms, which would serve as a solid foundation for these investigations.

For many non-uniformly hyperbolic systems beyond continued fractions, we do not have sharp error bounds for the Ulam approximation of Hausdorff dimension. The available estimates are often too pessimistic, and there is a gap between theory and observed numerical convergence. In higher dimensions (three or more), uniform grids become infeasible due to the curse of dimensionality. Sparse and tensor-train methods are promising, but their convergence properties for fractal sets are not yet fully understood. Real-world data is always noisy, and extending the framework to stochastic systems or to situations with measurement noise is an active area of research. Finally, while the transfer operator originated in classical dynamics, similar objects appear in quantum chaos and in general relativity. The idea of discretizing infinite-dimensional operators to extract spectral information may find applications in these fields as well.

Final Conclusions

The numerical path from Ulam's method to Hausdorff dimension is more than a computational recipe; it is a philosophical statement about the unity of mathematics. The geometry of a fractal, the statistics of a chaotic orbit, and the spectrum of a linear operator are revealed to be different facets of the same underlying dynamical reality. By discretizing the infinite, we bring this reality within reach of computation—and in doing so, we honor the spirit of exploration that Ulam embodied.

Funding: This research received no external funding.

Conflict of Interest: The author declares no conflict of interest.

References

1. Bowen, R. (1979). *Equilibrium States and the Ergodic Theory of Anosov Diffeomorphisms*. Springer Lecture Notes in Mathematics, Vol. 470.

2. Boyd, S. and Vandenberghe, L. (2020). *Introduction to Applied Linear Algebra: Vectors, Matrices, and Least Squares*. Cambridge University Press.

3. Dellnitz, M. and Junge, O. (1999). On the Approximation of Complicated Dynamical Behavior. *SIAM Journal on Numerical Analysis*, 36(2), 491–515.

4. Ding, J. and Zhou, A. (1993). Finite approximations of Frobenius–Perron operators. A solution of Ulam's conjecture to multi-dimensional transformations. *Physica D*, 92(1-2), 61–68.

5. Froyland, G. (2001). Extracting dynamical behavior via Markov models. In *Nonlinear Dynamics and Statistics*, Birkhäuser, 281–321.

6. Lasota, A. and Mackey, M. C. (1994). *Chaos, Fractals, and Noise: Stochastic Aspects of Dynamics*. Springer.

7. Ruelle, D. (1989). *The Thermodynamic Formalism*. Cambridge University Press.

8. Sebe, G. I. and Lascu, D. (2025). *Metrical and Ergodic Theory of Continued Fraction Algorithms*. Frontiers in Mathematics. Birkhäuser, Cham.

9. Ulam, S. M. (1960). *A Collection of Mathematical Problems*. Interscience Publishers.