Title
Chaos forgets and remembers: Measuring information creation, destruction, and storage

Permalink
https://escholarship.org/uc/item/3g49016g

Journal
Physics Letters A, 378(30-31)

ISSN
0375-9601

Authors
James, Ryan G
Burke, Korana
Crutchfield, James P

Publication Date
2014-06-01

DOI
10.1016/j.physleta.2014.05.014

Peer reviewed

Chaos forgets and remembers: Measuring information creation, destruction, and storage

Ryan G. James [a], Korana Burke [a], James P. Crutchfield [a,b,*]

[a] *Complexity Sciences Center and Physics Department, University of California at Davis, One Shields Avenue, Davis, CA 95616, USA*
[b] *Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA*

ARTICLE INFO

ABSTRACT

The hallmark of deterministic chaos is that it creates information—the rate being given by the Kolmogorov–Sinai metric entropy. Since its introduction half a century ago, the metric entropy has been used as a unitary quantity to measure a system's intrinsic unpredictability. Here, we show that it naturally decomposes into two structurally meaningful components: A portion of the created information—the ephemeral information—is forgotten and a portion—the bound information—is remembered. The bound information is a new kind of intrinsic computation that differs fundamentally from information creation: it measures the rate of active information storage. We show that it can be directly and accurately calculated via symbolic dynamics, revealing a hitherto unknown richness in how dynamical systems compute.

The world is replete with systems that generate information—information that is then encoded in a variety of ways: Erratic ant behavior eventually leads to intricate, structured colony nests [1,2]; thermally fluctuating magnetic spins form complex domain structures [3]; music weaves theme, form, and melody with surprise and innovation [4]. We now appreciate that the underlying dynamics in such systems is frequently deterministic chaos [5,6]. In others, the underlying dynamics appears to be fundamentally stochastic [7]. For continuous-state systems, at least, one operational distinction between deterministic chaos and stochasticity is found in whether or not information generation diverges with measurement resolution [8]. This result calls back to Kolmogorov's original use [9] of Shannon's mathematical theory of communication [10] to measure a system's rate of information generation in terms of the metric entropy. Since that time, metric entropy has been understood as a unitary quantity. Whether deterministic or stochastic, it is a system's degree of unpredictability. Here, we show that this is far too simple a picture—one that obscures much.

To ground this claim, consider two systems. The first, a fair coin: Each flip is independent of the others, leading to a simple uncorrelated randomness. As a result, no statistical fluctuation is predictively informative. For the second system consider a stock traded via a financial market: While its price is unpredictable, the direction and magnitude of fluctuations can hint at its future behavior. (This, at least, is the guiding assumption of the now-global financial engineering industry.) We make this distinction rigorous here, dividing a system's information generation into a component that is relevant to temporal structure and a component divorced from it. We show that the structural component captures the system's internal information processing and, therefore, is of practical interest when harnessing the chaotic nature of physical systems to build novel machines and devices [11]. We first introduce the new measures, describe how to interpret and calculate them, and then apply them via a generating partition to analyze several dynamical systems—the Logistic, Tent, and Lozi maps—revealing a previously hidden form of active information storage.

We observe these systems via an optimal measuring instrument—called a generating partition—that encodes all of their behaviors in a *stationary process*: A distribution $\Pr(\ldots, X_{-2}, X_{-1}, X_0, X_1, X_2, \ldots)$ over a bi-infinite sequence of random variables with shift-invariant statistics. A contiguous block of observations $X_{t:t+\ell}$ begins at index t and extends for length ℓ. (The index is inclusive on the left and exclusive on the right.) If an index is infinite, we leave it blank. So, a process is compactly denoted $\Pr(X_:)$. Our analysis splits $X_:$ into three segments: the *present* X_0, a single observation; the *past* $X_{:0}$, everything prior; and *future* $X_{1:}$, everything that follows.

* Corresponding author.
  E-mail addresses: rgjames@ucdavis.edu (R.G. James), kburke@ucdavis.edu (K. Burke), chaos@ucdavis.edu (J.P. Crutchfield).
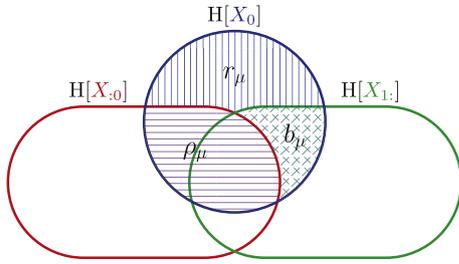
Fig. 1. A process's I-diagram showing how the past $X_{:0}$, present X_0, and future $X_{1:}$ partition each other into seven distinct information *atoms*. We focus only on the four regions contained in the *present information* $H[X_0]$ (blue circle). That is, the present decomposes into three components: ρ_μ (horizontal lines), r_μ (vertical lines), and b_μ (diagonal crosshatching). The *redundant* information ρ_μ overlaps with the *past* $H[X_{:0}]$; the *ephemeral* information r_μ falls outside both the *past* and the *future* $H[X_{1:}]$. The *bound* information b_μ is that part of $H[X_0]$ which is in the future yet not in the past. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The information-theoretic relationships between these three random variable segments are graphically expressed in a Venn-like diagram, known as an I-diagram [12]; see Fig. 1. The rate h_μ of information generation is the amount of new information in an observation X_0 given all the prior observations $X_{:0}$:

$$h_\mu = H[X_0|X_{:0}], \tag{1}$$

where $H[Y|Z]$ denotes the Shannon conditional entropy of random variable Y given variable Z. This quantity arises in various contexts and goes by many names: e.g., the Shannon entropy rate and the Kolmogorov–Sinai metric entropy, mentioned above [8]. The complement of the entropy rate is the *predicted information* ρ_μ:

$$\rho_\mu = I[X_{:0} : X_0], \tag{2}$$

where $I[Y : Z]$ denotes the mutual information between random variables Y and Z [12]. Hence, ρ_μ is the information in the present that can be predicted from prior observations. Together, we have a decomposition of the information contained in the present: $H[X_0] = h_\mu + \rho_\mu$.

A simple application of the entropy chain rule [12] to Eq. (1) leads us to a different view:

$$h_\mu = I[X_0 : X_{1:}|X_{:0}] + H[X_0|X_{:0}, X_{1:}]$$
$$= b_\mu + r_\mu. \tag{3}$$

This introduces two new information measures:

$$b_\mu = I[X_0 : X_{1:}|X_{:0}] \quad \text{and} \tag{4}$$

$$r_\mu = H[X_0|X_{:0}, X_{1:}]. \tag{5}$$

That is, created information (h_μ) decomposes into two parts: information (b_μ) shared by the present and the future but not in the past and information (r_μ) in the present but in neither the past nor the future.

The r_μ component was first studied by Verdú and Weissman [13] as the *erasure entropy* (their H^-) to measure information loss in erasure channels. To emphasize that it is information existing only in a single moment—created and then immediately forgotten—we refer to r_μ as the *ephemeral information*. The second component b_μ we call the *bound information* since it is information created in the present that the system stores and that goes on to affect the future.[1] It was first studied as a measure

of "interestingness" in computational musicology by Abdallah and Plumbley [14]. For a more complete analysis of this decomposition, as well as computation methods and related measures, see Ref. [15].

Isolating the information $H[X_0]$ contained in the present and identifying its components provides the partitioning illustrated in Fig. 1. This is a particularly intuitive way of thinking about the information contained in an observation. While some behavior (ρ_μ) can be predicted, the rest ($h_\mu = b_\mu + r_\mu$) cannot. Of that which cannot be predicted, some (b_μ) plays a role in the future behavior and some (r_μ) does not. As such, this is a natural decomposition of a time series; one that results in a semantic dissection of the entropy rate.

By way of an example, consider a few simple processes and how their present information decomposes into these three components. A periodic process of alternating 0s and 1s ($\ldots 01010101 \ldots$) has $H[X_0] = 1$ bit since 0s and 1s occur equally often. Given a prior observation, one can accurately predict exactly which symbol will occur next and so $H[X_0] = \rho_\mu = 1$ bit, while $r_\mu = b_\mu = 0$ bits. On the other extreme is a fair coin flip. Again, each outcome is equally likely and so $H[X_0] = 1$ bit. However, each flip is independent of all others and so $H[X_0] = r_\mu = 1$ bit, while $\rho_\mu = b_\mu = 0$ bits.

Between these two extrema lie interesting processes: those with *stochastic structure*. Processes expressing a fixed template, like the periodic process above, contain a finite amount of information. Those with stochastic structure, however, constantly generate information and store it in the form of patterns. Being neither purely predictable nor independently random, these patterns are captured by b_μ. The more intricate the organization, the larger b_μ. More to the point, generating these patterns requires intrinsic computation in a system—information creation, storage, and transformation [16]. We propose b_μ as a simple method of discovering this type of physical computation: Where there are intricate patterns, there is sophisticated processing.

How useful is the proposed decomposition and its measures? To answer this we analyze several discrete-time chaotic dynamical systems—the Logistic and Tent maps of the interval and the Lozi map of the plane—uncovering a number of novel properties embedded in these familiar and oft-studied systems. As an independent calibration for the measures, we employ Pesin's theorem [17]: h_μ is the sum of the positive Lyapunov characteristic exponents (LCEs). The maps here have at most one positive LCE λ, so $h_\mu = \max\{0, \lambda\}$. The symbols $s_0, s_1, s_2, \ldots, s_N$ for each process we analyze come from a generating partition. We produce a long sample of $N \approx 10^{10}$ symbols, extracting subsequence statistics via a sliding window.[2] Each window consists of a past, present, and future symbol sequence and we estimate r_μ and b_μ using truncated forms of Eqs. (4) and (5).

Consider first the Logistic map, perhaps one of the most studied chaotic systems:

$$x_{n+1} = a x_n (1 - x_n), \tag{6}$$

where $a \in [0, 4]$ is the control parameter and the initial condition is $x_0 \in [0, 1]$. Its generating partition is defined by:

$$s_n = \begin{cases} 0 & \text{if } x_n < \frac{1}{2}, \\ 1 & \text{if } x_n \geq \frac{1}{2}. \end{cases} \tag{7}$$

---

[1] Our terminology avoids the misleading use of the phrase "predictive information" for b_μ. The latter is not the amount of information needed to predict the future. Rather, it is part of the *predictable* information—that portion of the future which *can* be predicted.

[2] Window width is adaptively chosen in inverse proportion to the LCE. When the latter is low we use a longer window than when the system is fully chaotic. The minimum window width of $L = 31$ and adaptive widths were chosen so that numerical estimates varied by less than 0.01% when the width is incremented.
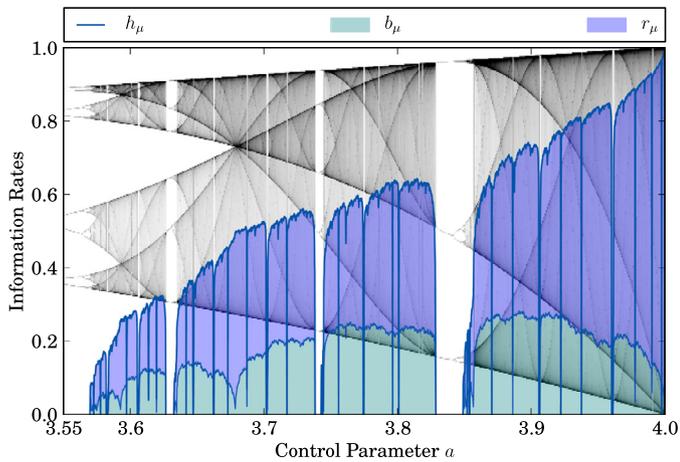
Fig. 2. Logistic map information anatomy as a function of control parameter a: Bound information b_μ is the lower (green shaded) component; ephemeral information r_μ is the upper (blue shaded) component. Entropy rate is the top (blue) line: $h_\mu = b_\mu + r_\mu$. As reference to the dynamical behavior, the map's bifurcation diagram is displayed in the background.

Fig. 2 shows the resulting measures as a function of control a, with the map's bifurcation diagram displayed in the background for reference.

The first point of interest is that the system's information generation is, in fact, a mixture of ephemeral (r_μ) and bound (b_μ) informations at nearly all chaotic ($h_\mu > 0$) parameter values. The second is that the division into the two components varies in a nontrivial way as a function of the control parameter a. Moreover, the boundary between the two appears nondifferentiable. At first blush, this is not surprising given that their sum $h_\mu (= \lambda)$ is known to be nondifferentiable. Finally, b_μ vanishes nontrivially only at parameters that coincide with the merging of the chaotic bands (e.g., $a = 4.0, 3.67857\ldots, 3.59257\ldots, \ldots$). Thus, the information generated by the Logistic map at these parameters is entirely forgotten.

Is the complex and nondifferentiable boundary between r_μ and b_μ simply a consequence of the entropy rate's complicated behavior or due a dynamical mechanism distinct from information creation? We answer this by analyzing the Tent map:

$$x_{n+1} = \frac{a}{2}\left(1 - 2\left|x_n - \frac{1}{2}\right|\right), \tag{8}$$

where $a \in [0, 2]$ is the control parameter. The generating partition for the Tent map is the same as for the Logistic map. Since the Tent map is piecewise linear, its Lyapunov exponent is simply $\lambda = \log_2 a$ and, by Pesin's theorem, so is the information generation $h_\mu = \log_2 a$; a rather smooth parameter dependence. As a result, the intricate structures exhibited in the Tent map's bifurcation diagram cannot be resolved by studying solely the behavior of the Lyapunov exponent (or h_μ) itself. Fig. 3 demonstrates that, despite the entropy rate's simple logarithmic dependence on control, its decomposition $h_\mu = b_\mu + r_\mu$ is not a smooth function of a. To emphasize, in sharp contrast with h_μ's simplicity, r_μ and b_μ again appear nondifferentiable—a complexity masked by the smooth h_μ. Thus, the two informational components capture a property in the chaotic system's behavior that is both quantitatively and qualitatively new. As with the Logistic map, we once again find that the bound information vanishes and that all of the information the Tent map generates is forgotten ($h_\mu = r_\mu$) at parameters corresponding to merging of chaotic bands ($a = 2^{1/2^k}$, $k = 0, 1, 2, \ldots$). In the Supplementary material we show how to calculate b_μ and r_μ in closed form for the Tent map at Misiurewicz parameters.
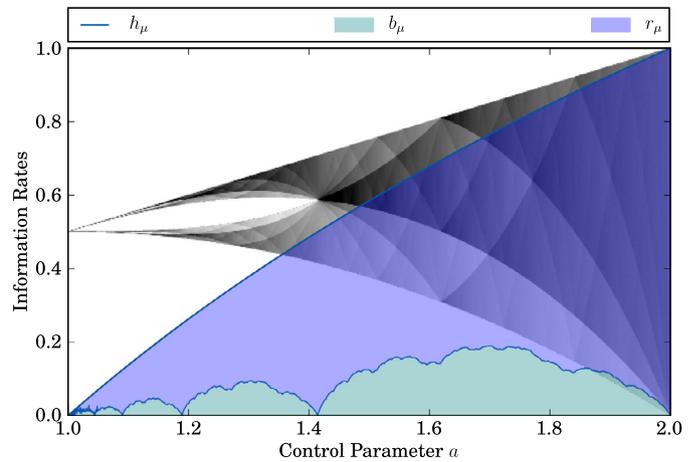
Fig. 3. Tent map information anatomy: Although $h_\mu = b_\mu + r_\mu$ is a smooth function of control—$h_\mu = \log_2 a$—the decomposition into bound and ephemeral informations is not. Graphics layout as in previous figure.
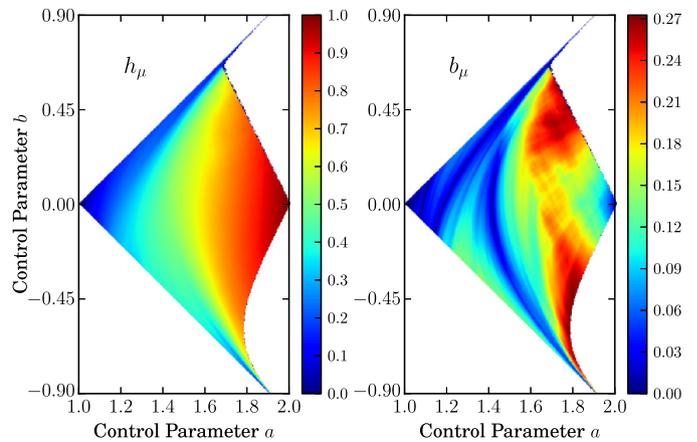
Fig. 4. Lozi map information anatomy: (Left) h_μ as a function of controls a and b. (Right) b_μ similarly. b_μ is maximized on the upper-right and lower-right edges of the a–b region that supports an attractor near the origin.

To explore how these measures apply more generally, we extend information anatomy to two dimensions by analyzing the Lozi map:

$$x_{n+1} = 1 - a|x_n| + y_n$$
$$y_{n+1} = bx_n. \tag{9}$$

The map exhibits an attractor near the origin within a diamond-shaped parameter region inside $(a, b) \in [1, 2] \times [-0.9, 0.9]$. Note that when $b = 0$ the map becomes isomorphic to the Tent map. The generating partition is given by:

$$s_n = \begin{cases} 0 & \text{if } x_n < 0, \\ 1 & \text{if } x_n \geq 0. \end{cases} \tag{10}$$

Fig. 4 shows h_μ (left) and b_μ (right) in the attracting parameter region. Mirroring the Tent map, the Lozi map's entropy rate varies smoothly over the attractor region, whereas b_μ varies in a more complicated manner. There are swaths of low b_μ corresponding to "fuzzy" mergings of chaotic bands. Notably, while the maximal h_μ occurs along the line $b = 0$, maximal b_μ occurs far from $b = 0$. Hence, large h_μ does not necessarily imply large bound information b_μ.

To sum up, we showed that a process's information creation rate decomposes, via a chain rule, into two structurally meaningful components. The components, the ephemeral information r_μ

and the bound information b_μ, provide direct insights into a system's behavior without detailed modeling or appealing to domain-specific knowledge. That is to say, they are relatively easily defined measures that can be straightforwardly estimated. More to the point, however, b_μ is a strong indicator of intrinsic computation. While related to information generation, we demonstrated that it captures a different kind of informational processing—a mechanism that actively stores information.

Concretely, decomposing information creation in the symbolic dynamics of the Logistic, Tent, and Lozi systems delineated the topography of their intrinsic-computation landscape. Awareness of this rich (and previously hidden) landscape will lead to improved engineering of natural systems as substrates for information processing [11]. And, it will lead to an expanded understanding of evolved information processing systems, such as the linguistic processes comprising human natural languages. A sequel will develop the decomposition further, including a geometric interpretation of active information storage that parallels the geometric view of information creation expressed in the Lyapunov exponents.

Acknowledgements

Appendix A. Supplementary material

Supplementary material related to this article can be found online at http://dx.doi.org/10.1016/j.physleta.2014.05.014.

References

[1] E. Bonabeau, G. Theraulaz, M. Dorigo (Eds.), Swarm Intelligence: From Natural to Artificial Systems, Oxford University Press, New York, 1999.
[2] S. Camazine, J.-L. Deneubourg, N.R. Franks, J. Sneyd, G. Theraulas, E. Bonabeau, Self-Organization in Biological Systems, Princeton University Press, New York, 2003.
[3] J.J. Binney, N.J. Dowrick, A.J. Fisher, M.E.J. Newman, The Theory of Critical Phenomena, Oxford University Press, Oxford, 1992.
[4] B.R. Simms, Music of the Twentieth Century: Style and Structure, second edition, Schirmer Books, New York, 1996.
[5] S.H. Strogatz, Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering, Westview Press, 1994.
[6] J.B. Jose, E.J. Saletan, Classical Dynamics: A Contemporary Approach, Cambridge University Press, New York, 1998.
[7] R.F. Streater, Statistical Dynamics: A Stochastic Approach to Nonequilibrium Thermodynamics, second edition, Imperial College Press, London, 2009.
[8] P. Gaspard, X.-J. Wang, Noise, chaos, and (ϵ, τ)-entropy per unit time, Phys. Rep. 235 (6) (1993) 291–343.
[9] A.N. Kolmogorov, Entropy per unit time as a metric invariant of automorphisms, Dokl. Akad. Nauk SSSR 124 (1959) 754 (in Russian); Math. Rev. 21 (2035b).
[10] C.E. Shannon, A mathematical theory of communication, Bell Syst. Tech. J. 27 (1948) 379–423, 623–656.
[11] W.L. Ditto, A. Miliotis, K. Murali, S. Sinha, M.L. Spano, Chaogates: morphing logic gates that exploit dynamical patterns, Chaos 20 (3) (2010) 037107.
[12] R.W. Yeung, Information Theory and Network Coding, Springer, New York, 2008.
[13] S. Verdú, T. Weissman, The information lost in erasures, IEEE Trans. Inf. Theory 54 (11) (2008) 5030–5058.
[14] S.A. Abdallah, M. Plumbley, Information dynamics: patterns of expectation and surprise in the perception of music, Connect. Sci. 21 (2) (June 2009) 89–117.
[15] R.G. James, C.J. Ellison, J.P. Crutchfield, Anatomy of a bit: information in a time series observation, Chaos 21 (3) (2011) 1–15.
[16] J.P. Crutchfield, K. Young, Inferring statistical complexity, Phys. Rev. Lett. 63 (1989) 105–108.
[17] Y.B. Pesin, Characteristic Lyapunov exponents and smooth ergodic theory, Russ. Math. Surv. 32 (4) (1977) 55–114.