Beyond the Digital Hegemony

A Focus Issue on
"Intrinsic and Designed Computation:
Information Processing in Dynamical Systems"

James P. Crutchfield,[1,2,*] William L. Ditto,[3,†] and Sudeshna Sinha[4,‡]

[1]*Complexity Sciences Center & Physics Department,*
University of California at Davis,
One Shields Avenue, Davis, CA 95616
[2]*Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501*
[3]*Harrington Department of Bioengineering*
Arizona State University, Tempe, Arizona 85287-9309
[4]*Indian Institute of Science Education and Research (IISER), Mohali,*
Transit Campus: MGSIPAP Complex, Sector 26, Chandigarh 160 019, India
(Dated: November 9, 2010)

How dynamical systems store and process information is a fundamental question that touches a remarkably wide set of contemporary issues: From the breakdown of Moore's scaling laws—that predicted the inexorable improvement in digital circuitry—to basic philosophical problems of pattern in the natural world. It's a question that also returns one to the earliest days of the foundations of dynamical systems theory, probability theory, mathematical logic, communication theory, and theoretical computer science. We introduce the broad and rather eclectic set of articles in this Focus Issue that highlight a range of current challenges in computing and dynamical systems.

Keywords: information processing, intrinsic computation, designed computation, dynamical systems, history

PACS numbers: 89.70.+c 05.45.-a 89.75.Kd

Contents

The reign of digital computing is being challenged, not only by fundamental physical limits but also by alternative information processing paradigms. The Focus Issue on Intrinsic and Designed Computation asks what the theory of nonlinear dynamical systems has to offer in response. Historical reflection on the origins of information and computation theories reveals the formerly close connections to dynamical systems and, in particular, to the first concrete ways to measure deterministic chaos. The articles in the collection, intentionally, vary quite widely in their views on these issues, from the most abstract formal settings, in which determining the very chaoticity of a dynamical system appears to be as hard as solving the hardest mathematical problems, to the most concrete *in silico* implementations of chaotic logic. The technological promise is substantial: faster, less expensive, and more energy efficient computing. Perhaps the most long-lasting impact, though, will be a new appreciation of the ubiquity of information processing in the natural world.

I. INTRODUCTION

Many contemporary research domains face, at their core and perhaps unawares, the confounding problems of defining and measuring information processing in dy-

———

[*]Electronic address: chaos@cse.ucdavis.edu
[†]Electronic address: william.ditto@asu.edu
[‡]Electronic address: sudeshna@imsc.res.in (on leave from: The Institute of Mathematical Sciences, Chennai, India)

namical systems. These range from technology to fundamental science and, even, epistemology of science:

The 2020 Digital Roadblock: The arrested acceleration of microprocessor computing power over the last decade gives evidence to the failure, after half a century of success, of Gordon Moore's scaling laws for digital technology [1–3]. The move to multiple CPU cores does not solve the basic problem: By 2020 the speed, spatial scale, circuit size, energy dissipation, and manufacturing economics required by Moore's law are not sustainable with current solid-state physics. One solution path is to call for a radical rethinking of what it means for natural and engineered systems to store and process information. What natural substrates are complex enough to support controllable and useful information processing? New substrates promise an end-run around the 2020 roadblock. But how will we identify and evaluate proposed substrates for their ability to support information storage and logic? Developing principled (and practical) measures of information processing will go some distance to helping solve this problem. Appreciating how the richness of nonlinear dynamics supports information processing will go some distance to radically new substrates for computing.

The Central Dogma of Neurobiology: Neuroscience studies cell types, tissues, and organs that ostensibly evolved to store, transmit, and process information. That is, the behavior and organization of neural systems support computation in the service of adaptation and intelligence. How are the intricate physical, biochemical, and biological components structured and coordinated to support natural, intrinsic neural computing?

Physical Intelligence: Does intelligence require biology, though? Or can there be alternative non-biological substrates which support system behaviors that are to some degree "smart". The possibility of abiotic intelligence was an underpinning of artificial intelligence [4]. Generally, the successes in that research paradigm, begun well over half a century ago, have been severely limited. And those limitations call into question its historical origins in mathematical logic.

More to the point, is intelligence a subset, a particular kind, of discrete symbolic computation? On reflection, this view might very well have the relation inverted and so be responsible for AI's difficulties. That is, perhaps digital computation is one form of intelligence embedded in a noisy, continuous, distributed natural world. Are there other approaches to intelligence? Can intelligence be separated from the physical world also? Or does it have an inextricably physical, non-logic quality?

To be concrete, what would a physically intelligent abiotic system look like? Given such a thing, how would one detect its intelligence? Can we move beyond the subjectivity of the Turing Test—"We know it when we see it"—to physics-based criteria—"We know it when we measure it"?

Emergent Patterns: Intelligence aside, how do we define and detect spontaneous organization, in the first place? Can we obviate the effects of observer-dependence when detecting emergent properties? That is, can we do good and consistent science when considering nature's patterns? These are key and common problems that confront the researcher probing any new complex system.

II. INTRINSIC COMPUTATION

The question, then, is how to address such challenges, leveraging the substantial inheritances from statistical physics, nonlinear dynamics, and information and computation theories. The premise of intrinsic computation is that there is a concrete relationship between pattern and information processing. Intrinsic computation attempts to answer three question about any dynamical system:

1. How much historical information does a process store?

2. In what architecture is that information stored?

3. How is that stored information transformed to produce future behavior and organization?

III. DESIGNED COMPUTATION

Whether one agrees with these three particular questions, as identifying a kind of computation in dynamical systems, most would agree that the computing machines we build today do something much more. They process information in ways that are *useful* to us. Intrinsic computation makes no reference to utility. Designed computation, unless part of an art installation, say, is nothing if not useful.

There is a physics question in all of this, what is the relationship between intrinsic and designed computation? They would seem to be related. Is intrinsic computation necessary? That is, to do n bits of useful computation requires at least n bits of intrinsic computation in the supporting substrate. Certainly, intrinsic computation is not sufficient. We waste vast amounts of physical degrees of freedom and their attendant intrinsic computation in our current computing implementations. Memory chips sit idle, while the CPU silicon actively burns a furious amount of energy.

More than a performance gap, though, there is a semiotic gap between intrinsic and useful computation. How does utility factor into the fundamental physical limits of computation? Can it? In the difference between intrinsic and designed computation lie many challenges. On the one hand, the questions of the semiotic gap are immediate: They play out in our palms and on our laps and desktops everyday. On the other, constructive answers could well break the prevailing digital technology lock-in and, as a by-product, enrich the sciences of the information age.

IV. THE LARGER HISTORICAL SETTING

The search to quantify how the behavior of a dynamical system supports or constrains information processing has a long and venerable history. One of the first fruitful results came in the 1950s with the Soviet mathematician Andrei Kolmogorov's adaptation of Claude Shannon's concept of information [5–7] to measure the degree of unpredictability of nonlinear dynamical systems [8–11]. This led, shortly thereafter, to Ornstein's solution [12, 13] to the famous isomorphism problem [28]: When are two dynamical systems equivalent? This work built a solid bridge between the study of chaotic dynamical systems and information theory. It gave, in principle, a procedure for distinguishing differently organized complex systems. A few short years later, probing the algorithmic origins of randomness, Kolmogorov and Gregory Chaitin (and others) [11, 14–18] recast the framework in terms of Alan Turing's theory of computation [19]. This demonstrated how a complex system's behavior can be construed as a computation. Ever since these transdisciplinary innovations, dynamical systems, information theory, and computation theory have been key tools for probing how complex systems store and generate information—that is, how they compute intrinsically.

In many ways, Kolmogorov's pioneering connections were anticipated, following as they did the efflorescence of *cybernetics* during the 1940s and 1950s [20, 21]. Norbert Wiener, coiner of the label and central protagonist, focused on control and communication in biological and engineered systems. Despite being a consummate mathematician, his style of cybernetics foundered on the sheer analytical difficulty of extending informational measures to continuous-valued processes. In practical historical terms, cybernetics was eclipsed by the successes of Shannon's mathematical theory of communication [22]. This was partly due to Shannon's technically adroit concentration on discrete-valued processes and partly to his showing the way to error-free coding which gave information theory much traction in engineering communication systems. Nonetheless, Wiener's eloquent vision stands today as a challenge to modern dynamical systems and how their information processing arises in natural systems and can be used to build new generations of computing device [23, 24]. Many of the mathematical problems that stymied him half a century ago, however, remain as roadblocks today to the theory and application of dynamical systems and statistical mechanics.

A similarly prescient vision for complex systems and, in particular, for identifying and measuring complexity, was given by Warren Weaver in his 1948 essay on "Science and Complexity" [25]. There, he appears to have been the first to clearly distinguish randomness (his "disordered complexity") from structure (his "organized complexity"). The challenge Weaver leaves us even now is the central (and constructive) role that understanding complex systems must play in the sciences. Notably, its role in explaining social and cultural evolution had been anticipated two decades earlier by the philosopher Alfred North Whitehead [26]. There are also hints of this kind of "structural" thinking in the molecular basis of biology put forth by Erwin Schrödinger's in his well known 1944 Cambridge lecture *What is Life?* [27].

Going from the technical advances of Kolmogorov to the visions presented by Wiener, Weaver, and Whitehead, these examples serve to demonstrate the robust intellectual history of information processing in dynamical systems—how to define, detect, and harness it. Importantly for today, they reveal a long-lived, historical momentum that underpins much of the modern dynamical systems approach to science and engineering.

Perhaps, then, it is somewhat surprising to realize that many of the foundational issues raised in these earlier periods have not been solved, by any means. For example, while "disorganized complexity" is well articulated in the statistical mechanical origins of thermodynamic entropy and temperature and in the algorithmic foundations of randomness, there are still active debates on candidates for Weaver's "organized complexity". Moreover, much of the original enthusiasm and the predictions of such luminaries as Turing, Wiener, and John von Neumann for the automation of human intelligence and the control of nature have simply not been realized. Indeed, the rise of dynamical systems over the last several decades is testimony to how much more there is to understand and how many more applications are possible.

Fortunately, a new level of rigor, in concepts and in analysis, is now apparent in how statistical mechanics, information theory, and computation theory can be applied to dynamical systems. The meteoric rise of both computer power and machine learning has led to new algorithms that address many of the computational difficulties in managing data from complex systems and in estimating various information processing measures. Given progress on all these fronts, the time is ripe to develop a much closer connection between fundamental theory and applications in the many areas that consider intrinsic and designed computation.

V. FACILITATING THE FUTURE

Perhaps more important than a necessary amount of stock-taking, the Focus Issue presents the opportunity to push toward real progress on the original challenges—progress that could demonstrate to the larger scientific and engineering communities the benefits, in concrete quantitative terms, of the dynamical systems view of intrinsic and designed computation.

Naturally, one goal is to stimulate new thinking about what information and computation are. In looking far enough ahead, the Focus Issue offers solutions—that is, does more than present dynamical systems as a handmaiden to current technology applications. One might ask what new there is to do, especially in light of what was claimed: recent substantial progress. One way to il-

lustrate the potential is to re-examine how Shannon's original mathematical theory of communication [5] is used in nonlinear dynamics. It's helpful to recall that the theory actually consists of two components: The first, what one might call *information theory proper*, addresses what information is and how to quantify it. The second, *communication theory*, concerns how to efficiently and accurately transmit it.

Reviewing the last two decades' efforts to define and measure information processing, a large majority appeal only to information theory proper as a starting point to introduce this or that informational measure. The flip side is that the rich set of concepts and methods comprising communication theory—channels, codes, rate distortion theory, and the like—are greatly underutilized. The result is that there has been relatively little progress in analyzing the architecture of how complex systems support the flow and storage of information. A similar statement about the use of Shannon's informational approach to cryptographic systems [6] is even more true.

Despite being developed more that a half century ago, Shannon's communication and cryptographic theories provide insights essential today for a deeper understanding of information processing in natural and engineered systems. Thus, one rather direct strategy for analyzing the gap between intrinsic and designed computation is to explore the use of these theories in much more depth. Recalling the historical relationship between Shannon's success with discrete-valued processes and the difficulties of Wiener's cybernetics of continuous systems, the direct strategy for progress can be easily pushed further. One straightforward goal is to reconsider Wiener's original vision for cybernetics. One trusts that there has been sufficient technical progress in the intervening half century that at least several of cybernetics' original challenges will soon be surmounted.

VI. COMPUTING IN AND WITH DYNAMICAL SYSTEMS

CHAOS's Focus Issue brings together researchers from a variety of fields who consider intrinsic and designed computation from the perspectives of dynamical systems and statistical mechanics. Some of the questions addressed are:

1. What is the role of nonlinearity in computation?
2. Are there fundamental measures of information processing that can be applied across different physical, chemical, biological, and social systems?
3. How is a system's causal organization, reflected in models of its dynamics, related to its computational capability?
4. Can we reach agreement on general properties that all measures of information processing and computation must have?
5. How would the scientific and engineering communities benefit from a consensus on the properties that

measures of information processing should possess? One can imagine international standards for measuring the intrinsic computational capability—a machine IQ—of candidate substrates.

For all the reasons just outlined, the time was ripe to address these head-on. And this, in turn, led rather directly to the Focus Issue on Intrinsic and Engineered Computation. The result is a highly interdisciplinary group of contributors representing engineering, physics, chemistry, biology, neuroscience, computer science, and mathematics. An important goal was to understand the successes and difficulties in deploying these concepts in practice. And so, contributors from both theory and experiment are represented, with a particular emphasis on those that have constructively bridged the two. Here, now, is a brief preview of those contributions.

Nonlinear Semiconductor Lasers and Amplifiers for All-Optical Information Processing by Michael Adams, Antonio Hurtado, Dmitry Labukhin, and Ian Henning: The authors demonstrate the practical use of the nonlinear properties of semiconductor lasers and laser amplifiers for optical logic and chaotic communication.

The Complexity of Proving Chaoticity and the Church-Turing Thesis, by Cristian S. Calude, Elena Calude, and Karl Svozil: The authors argue that proving a dynamical system is chaotic is equivalent to solving the hardest problems in mathematics. They suggest, provocatively, that classical physical systems may compute the hard or even the uncomputable.

Numerical Information Processing under the Global Rule Expressed by the Euler-Riemann ζ Function Defined in the Complex Plane by Francoise Chatelin: The author describes how objects from number theory—variations on the famous ζ function—are deeply connected with numerical information processing. Her analysis gives new insight into a time-honored mathematical challenge—the critical line of the ζ function—and suggests a cognitive view of the Fourier transform.

Synchronization and Control in Intrinsic and Designed Computation: An Information-Theoretic Analysis of Competing Models of Stochastic Computation by James Crutchfield, Christopher Ellison, Ryan James, and John Mahoney: The authors describe the process of synchronization to a nonlinear dynamical system from an information-theoretic point of view. Building on past developments in the computational mechanics of ϵ-machines and related information measures, they introduce a new set of measures: the block-state and state-block entropies that allow one to analyze the convergence to synchronization. They introduce a new information-theoretic classification for finite-memory stochastic processes based on synchronization and controllability.

Distribution and Regulation of Stochasticity and Plasticity in Saccharomyces Cerevisiae by Roy Dar, David Karig, John Cooke, Chris Cox, and Michael Simpson: The authors address the important issue of fluctuations in nanoscale complex systems, analyzing a biological cell as a prototype for addressing the trade-off

between stochasticity and determinism. The gene expression of the budding yeast (*Saccharomyces cerevisiae*) under many different conditions reveals the balance between the deterministic and the stochastic response of genes when external stimuli change. The ideas presented should help to design better nanoscale devices where, due to the small size, the role of fluctuations differs substantially compared to bulk systems.

Chaogates: Morphing Logic Gates Designed to Exploit Dynamical Patterns by William Ditto, Aris Miliotis, Krishnamurthy Murali, Sudeshna Sinha, and Mark Spano: The authors show how to commandeer the wide variety of patterns that chaotic systems generate to implement the full panoply of logic gates. They describe how to design a dynamical computing device—the *chaogate*—that can be rapidly re-purposed to become any desired logic gate. In addition to reviewing the basic design principles, they extend the formalism to include asymmetric logic functions.

Intrinsic Information Carriers in Combinatorial Dynamical Systems by Russ Harmer, Vincent Danos, Jerome Feret, Jean Krivine, and Walter Fontana: The authors show how protein modularity, which derives from interaction specificity, underlies the vast combinatorial complexity characteristic of many biological networks. They argue that this effectively prevents their study via kinetic equations. Given that combinatorial complexity cannot be eliminated, they suggest capturing the system's average or deterministic behavior using a graph-based framework of rewrite rules—a new class of computational model—in which each rule specifies only the information that an interaction mechanism depends on. They demonstrate how to find aggregated variables that reflect the causal structure laid down by the mechanisms expressed by the rules—what they call *fragments*. Fragments are self-consistent descriptors of system dynamics whose time evolution is governed by a closed system of kinetic equations and which they go on to identify as the primary information carriers.

Information Modification and Particle Collisions in Distributed Computation by Joseph Lizier, Mikhail Prokopenko, and Albert Zomaya: The information dynamics of distributed computation is a topic of continuing interest in complex systems. The authors contribute to this subject by considering quantification of operations like information storage, transfer, and modification. In order to describe the dynamics of information in computation, the authors attempt to quantify these operations on a local scale in space and time, exploring how information modification can be quantified at each spatiotemporal point in a system. The techniques are tested on cellular automata, but the results are expected to be of wider application.

Discrete Analogue Computing with Rotor-Routers by James Propp: The author describes an intriguing computational model "rotor-routing"—tokens flow through a network—that implements asynchronous and distributed computation. Rotor-router networks are both discrete

analogues of continuous linear systems and deterministic analogues of stochastic processes. Propp shows that one can efficiently check that a rotor-router computation has been carried out correctly in less time than the computation itself required.

Optimal Causal Inference: Estimating Stored Information and Approximating Causal Architecture by Susanne Still, James Crutchfield, and Christopher Ellison: The authors introduce a statistical-mechanics approach to inferring the causal architecture of a stochastic dynamical system. They introduce two methods of causal inference: optimal causal filtering and optimal causal estimation. Through this new study, and through computational mechanics more generally, causal states have been demonstrated to be powerful representation for analyzing emergent organization in complex systems. The authors provide a quantitative way of exploring fundamental tradeoffs associated with model complexity versus model predictability. This is relevant to understanding how prediction error scales when one reduces a model and how these tradeoffs differ for differently structured processes.

Computing Adaptive Bayesian Inference from Multiple Sources by Hava Siegelmann and Lars Holzman: The authors discuss one of brain's most basic functions, namely integrating sensory data from diverse sources. This ability underscores a fundamental question: Is the neural system computationally capable of intelligently integrating data? This work gives a specific computational algorithm to confirm the view that appropriate neural architectures are able to calculate posterior probabilities and to learn relative reliabilities. The novelty in the paper is that the two types of calculations can be executed with a single algorithm.

How Does a Choice of Markov Partition Affect the Resultant Symbolic Dynamics? by Hiroshi Teramoto and Tamiki Komatsuzaki: The authors consider a question of abiding interest in the application of symbolic dynamics and, in fact, in any discrete measurement of a continuous-valued system. They show that, if a Markov partition is a map-refinement of another Markov partition, then one can uniquely translate the symbolic sequences from one Markov partition to those produced via the other or vice versa.

Nature Computes—Quantifying Information Processing in Quantum Dynamical Systems by Karoline Wiesner: The author reviews the current thinking on intrinsic quantum computation as a way to quantify the information processing in natural quantum systems, combining tools from dynamical systems theory, information theory, quantum mechanics, and computation theory. As one result, she gives upper and lower bounds for intrinsic computation of a quantum dynamical system.

VII. CLOSING REMARKS

The Focus Issue is just a beginning. To be fair, though, it is a continuation of long-lived research program that goes back the origins of dynamical systems over a century ago and to the first days of cybernetics, a half century ago. We will judge it a success if the contributions stimulate new thinking—new ways to go beyond the digital hegemony to radically new kinds of computing and a deeper appreciation of the role of natural and engineering information processing.

Acknowledgments

[1] G. E. Moore. Cramming more components onto integrated circuits. *Electronics*, 38(8):56–59, 1965.

[2] G. E. Moore. Progress in digital integrated electronics. *Technical Digest 1975. International Electron Devices Meeting, IEEE*, pages 11–13, 1975.

[3] G. E. Moore. Lithography and the future of moore's law. *Proceedings of the SPIE*, 2437:1–8, May 1995.

[4] P. McCorduck. *Machines Who Think: A Personal Inquiry into the History and Prospects of Artificial Intelligence*. A.K. Peters, Natick, Massachusetts, second edition, 2004.

[5] C. E. Shannon. A mathematical theory of communication. *Bell Sys. Tech. J.*, 27:379–423, 623–656, 1948.

[6] C. E. Shannon. Communication theory of secrecy systems. *Bell Sys. Tech. J.*, 28:656–715, 1949.

[7] C. E. Shannon. Prediction and entropy of printed English. *Bell Sys. Tech. J.*, 30:50, 1951.

[8] A. N. Kolmogorov. A new metric invariant of transient dynamical systems and automorphisms in Lebesgue spaces. *Dokl. Akad. Nauk. SSSR*, 119:861, 1958. (Russian) Math. Rev. vol. 21, no. 2035a.

[9] A. N. Kolmogorov. Entropy per unit time as a metric invariant of automorphisms. *Dokl. Akad. Nauk. SSSR*, 124:754, 1959. (Russian) Math. Rev. vol. 21, no. 2035b.

[10] Ja. G. Sinai. On the notion of entropy of a dynamical system. *Dokl. Akad. Nauk. SSSR*, 124:768, 1959.

[11] M. Li and P. M. B. Vitanyi. *An Introduction to Kolmogorov Complexity and its Applications*. Springer-Verlag, New York, 1993.

[12] D. S. Ornstein. *Ergodic Theory, Randomness, and Dynamical Systems*. Yale University Press, New Haven, 1974.

[13] D. S. Ornstein. Ergodic theory, randomness, and chaos. *Science*, 243:182, 1989.

[14] R. J. Solomonoff. A formal theory of inductive control. *Info. Control*, 7:224, 1964.

[15] A. N. Kolmogorov. Three approaches to the concept of the amount of information. *Prob. Info. Trans.*, 1:1, 1965.

[16] G. Chaitin. On the length of programs for computing finite binary sequences. *J. ACM*, 13:145, 1966.

[17] P. Martin-Lof. The definition of random sequences. *Info. Control*, 9:602, 1966.

[18] A. N. Kolmogorov. Combinatorial foundations of information theory and the calculus of probabilities. *Russ. Math. Surveys*, 38:29, 1983.

[19] A. M. Turing. On computable numbers, with an application to the entsheidungsproblem. *Proc. Lond. Math. Soc. Ser. 2*, 42:230, 1936.

[20] N. Wiener. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series*. Wiley, New York, 1949.

[21] N. Wiener. Nonlinear prediction and dynamics. In *Norbert Wiener, Collected Works III*. MIT Press, Cambridge, 1981.

[22] F. Conway and J. Siegelman. *Dark Hero of the Information Age: In Search Of Norbert Wiener—Father of Cybernetics*. Basic Books, New York, 2004.

[23] N. Wiener. *Cybernetics: Or Control and Communication in the Animal and the Machine*. MIT, Cambridge, 1948.

[24] N. Wiener. *The Human Use Of Human Beings: Cybernetics And Society*. Da Capo Press, Cambridge, 1988.

[25] W. Weaver. Science and complexity. *American Scientist*, 36:536, 1948.

[26] A. N. Whitehead. *Process and Reality*. The Free Press, New York, corrected edition, 1978.

[27] E. Schrodinger. *What is Life? The Physical Aspect of the Living Cell*. Cambridge University Press, Cambridge, United Kingdom, 1944.

[28] The isomorphism problem was part of the constructivist program in mathematics laid out by David Hilbert in 1900.