The Thermodynamic Origin of Structure

William Rogers*

April 5, 2019

Abstract

James Crutchfield and Co. have developed a framework they call computational mechanics which presents a mathematical definition of structure as how much of a system's past or future states can be ascertained given its present state. I hypothesize that structure arises in systems always and only when energy is dissipated at a higher rate than a system can accommodate by diffusing the energy evenly in all directions. I call this *structured dissipation.* The model system used to illustrate structured dissipation is the transition from conduction, an unstructured process, to convection, a structured process, that occurs when a liquid system attached to a heat sink is heated above a critical point. The paper concludes with thoughts about how to test this theory through the examination of the emergence of structure in other physical systems.

1 Introduction

Despite structure being precisely what makes the world nonrandom, nontrivial, and interesting, there is no general theory of the behavior of structure. There are plenty of theories that exist to explain the particular behavior of particular structures–this is the content of most science–but so far it has not been possible to find basic commonalities between all manifestations of structure that would enable constructing a theory of structure from fundamental principles.

In large part, the behavior of structure has eluded theorization because structure itself has remained undefined in any precise or measurable sense. Commonly, structure is thought of as the relations between parts of a system. What is totally uniform or totally random has no structure because there are no particular relations or meaningful correlations between the parts. "Relations" can take many different forms in systems, which makes it challenging to posit a universal metric of measurement for such relations.

James Crutchfield and company have, over the last 30 years, developed a theory that attempts to universally identify and measure structure. Crutchfield's theory of computational mechanics claims that a system has structure when information about past and future states of the system can be learned from observing the present state of the system. The more structure a system has, the more information a system "stores" in its present state about

*The author can be reached at williamrogersthe3rd@gmail.com.

future and past states. Connecting this definition to the common notion of structure noted in the previous paragraph, the amount of structure in a system is the amount of internal relations, which can be measured by the predictability of the system's past and future states given its present state. In order to build this theory up from first principles, it has hitherto been applied mostly to extremely simple abstract and real-world systems.

What is key to this definition is the way it delineates structure from order or uniformity, which in the past have often been conflated with structure. A totally ordered system is one that is entirely uniform in its progression between states, e.g. it moves through state "A", then goes to state "A" again, and so on, which can be represented as "AAA...". Whereas structure in a process implies that the system passes through different states in such a way that looking at one state, one can predict future and past states. For example, a system that progresses through states "ABCABCABC..." is structured because if you see the system in state "B", one knows that previously it was in state "A" and subsequently it will be in state "C".

On the opposite side of the spectrum from order lies disorder, randomness, and entropy, terms which are used interchangeably in this paper. A system that randomly moves between states, like "AABACCBCABC...", is not structured because there is nothing in the present state that helps predict past or future states. Just as order has sometimes been mistaken for structure, so has disorder. For example, a very disordered dataset may appear like it has "a lot going on" and thus be the result of complex processes and require complex explanation. Distinguishing random phenomena from very complex phenomena may be tricky sometimes, but the distinction is critical: mistaking a random sequence of letters generated by, say, flipping a coin, for a subtly intricate system will generate a lot of wasted effort. A major finding of computational mechanics is that systems can have independently varying degrees of structure and order.

Crutchfield's definition of structure has the further of being mathematically formalized, such that the structure and entropy of systems can be quantified given an appropriate model of the system. The difficult part is modeling systems in such a way that they can be represented as having a linear sequence of states upon which the framework of structure can then be applied. Nonetheless, Crutchfield and colleagues have successfully applied their framework to a wide range of physical systems, from chaotic crystal formation to turbulence [1].

What is most titillating about this theory of structure is the possibility that structure is indeed a fundamental property of physical systems and not merely a human construct. Perhaps structure in the world is as real as entropy in a gas. If so, the investigation of structure today might be capable of opening up new domains of science, the way Newton and his contemporaries opened the floodgates with the formalization of the concept of energy. Furthermore, if structure is a fundamental property of systems, akin to energy or entropy, then how structure changes in systems over time might be governed by distinct laws akin to the laws of thermodynamics. One such law is proposed below as the thesis of this paper.

The paper proceeds as follows: In Section 2, I present an overview of Crutchfield's framework for defining structure. In Section 3, I argue that Crutchfield's concept of structure has deeper theoretical implications–that the behavior of structure in systems might be governed by universal laws like those in thermodynamics–and that an approach to structure in this vein is justified even if speculative. In Section 4, I illustrate how the transition from conduction to convection in a heated liquid is an instance of the emergence of structure. In Section 5, I generalize the convecting liquid example to argue that it is a universal law that all structure is the result of energy dissipation in a system above a particular threshold beyond which that energy is no longer capable of being dissipated uniformly in space outward from the source. The emergence of structure facilitates faster energy dissipation by increasing the overall channel capacity through which a system dissipates energy. In Section 6, I discuss how this theory might be further investigated and tested against the emergence of structure in other well-known physical systems.

2 Structure

A common notion of structure is that it is the relations between parts of a system [2, 3], and contemporary definitions of structure in physics attempt to quantify these relations. The more structure a system or process has, the more relations it has between its parts, and we describe such systems as complex. Other closely related and often vaguely-defined concepts are pattern, organization, and design.

2.1 Entropy

Before defining structure formally, it's useful to recognize what structure is not. A gas at equilibrium has almost no informative relations between particles because each one behaves independently of the rest, and the position of a gas particle at any time is entirely unrelated to the positions of the other particles. Thus, such a gas lacks structure. Conversely, a perfect crystal also has no meaningful relations because each particle lies adjacent to each other in the most rigid and least relationally complex way and so lacks structure[1].

While a gas in equilibrium and a perfect crystal each have no structure, an evenly distributed gas has a very high value of entropy while a perfect crystal has a very low value of entropy. Entropy here is meant in the sense Boltzmann used it, as the number of ways the micro-parts of a system can be arranged given some macro-level constraints. In colloquial language, this concept of entropy can be loosely approximated as the amount of disorder, or the expected degree of randomness in the arrangement of parts, in a system.

One can imagine analogs of these physical systems in the form of digital bit strings of 0s and 1s. The analog of a gas is a series of coin flips, where 1 signifies heads and 0 signifies tails. Even though the same coin may be flipped each time, the outcome of the last flip has no relation to the outcome of the next flip or 100th flip, just as the location of one particular gas particle has no relation to the particular location of any other gas particle at any particular time. A string of all 0s is the binary digit string analog of a crystal because each digit is totally ordered in a way that obviates any informative relations between them.

[1] The concept of relation can be interpreted generally, such that 'the same as' or 'randomly' could be said to be relations, but I am arguing that these are not relations in the sense I mean. I use the concept of relation here to mean *informative* relations, and this kind of relation necessarily implies that having knowledge about such a relation and a part of the system provides new information about other parts of the system in a way that 'the same as' and 'randomly' do not. This meaning becomes clearer once the framework below is fully introduced

Entropy in digital form was first quantified and formalized by Claude Shannon in the 1940s [4]. Shannon quantified information entropy as the degree of surprise or uncertainty contained in each new piece of information, which is also the same as the amount of randomness or disorder contained in a system between its parts[2]. The amount of information entropy in a process or bit string can be determined using the following equation:

$$H(X) = -\sum_{i=1}^{n} Pr(x_i) log_b Pr(x_i)$$

Where X represents a discrete variable that can take values from an alphabet of possible symbols $\{x_1, ..., x_n\}$, each with probability $Pr(x_i)$. The log base b can be set to 2 or e or to the length of the alphabet of possible symbols, and the choice of base changes the units that the equation outputs. For example, log base 2 outputs *bits*, which is equivalent to the number of yes/no questions needed to be answered to erase all the uncertainty and randomness in the object of study. Log base 2 is used in this paper for ease of comparison across different contexts.

Notice that entropy $H(X)$ is maximized when each symbol has equal probability such that there is the greatest uncertainty of what the next symbol will be—as when there are two symbols, heads and tails, and their probabilities are both 0.5, or when each symbol on a six-sided die has probability 0.17. Entropy is minimized when there is one possible symbol with probability 1 for each digit, as in a string of all 0s. The entropy of a coin-flip system can be varied continuously between maximum and minimum values by sliding the symbol probabilities between the extremes of equiprobability and a completely determined probability of 1[3]. These intermediate degrees of entropy can be thought of as a biased coin, where there might be more or less surprise from one flip to the next, but each flip still also bears no information about the outcomes of future or past flips and so there are no relations or structure in the system.

2.2 A quantitative definition of structure

While there's no consensus among physicists about how to define structure, the framework developed by James Crutchfield and colleagues has received the most attention[5]. Crutchfield calls his project computational mechanics—a kind of statistical mechanics for how

[2] Just as Shannon's concept of information entropy quantifies the degree of uncertainty or disorder given a process in digital form, the concept of thermodynamic entropy was re-framed in physical form by Boltzmann and Gibbs in the late 1800's as the number of microstates compatible with a given equilibrium macrostate, which laid the foundation for statistical mechanics. The underlying equations for Boltzmann's and Shannon's entropy are identical even though they nominally apply to different kinds of systems, physical and informational. The exact relation between these two measures is a topic of much research but the basic idea that they measure fundamentally the same quantity is axiomatic in much contemporary thermodynamics research. For more in-depth analysis, see [6, 7].

[3] The biased coin example clearly invalidates [8] the old notion, still held by some [9, 10], that structure and complexity are defined as lying mid-way between maximum and minimum entropy. On the contrary, a coin of any bias can be conceived that produces any intermediate value of entropy while still containing no structure. Structured objects sometimes do have intermediate values of entropy, but it is not that that makes them structured.

processes are structured[4].

Central to this framework is the mode of representation used to model systems which Crutchfield calls the ϵ-machine. The ϵ-machine is a map of a system's possible states, the transition probabilities between the states, and the output values of the system upon transitioning between states. The states of the ϵ-machine are what Crutchfield calls causal states, which partition the state space of a system into those that have unique future state probability distributions. Thus, those states that have the same causal relation to future states are all grouped together because they are causally equivalent. For example, in a coin flip system, the system is always in the same one causal state in relation to future states of the system because whatever the current state, all future outputs and states are equally likely. Crutchfield and colleagues have derived proofs that for any system, the ϵ-machine or an equivalent representation is both the minimal representation of that system and the optimally predictive one [11].

The ϵ-machine facilitates the calculation of two new quantities of a system: its entropy rate and its statistical complexity. The entropy rate is the information entropy per symbol that the system produces:

$$h_\mu = -\sum_{\sigma \in S} Pr(\sigma) \sum_{\{x\}} Pr(x|\sigma) log_b Pr(x|\sigma)$$

Where $Pr(\sigma)$ is the probability the system is in a causal state σ within the alphabet of all possible symbols S at any one time, which when calculated over each causal state gives the causal state probability distribution of the system. $Pr(x|\sigma)$ is the probability of the system transitioning from a state σ given measurement x, which when summed up produces the total transition probability distribution of the system. The μ in h_μ is meant to remind us that this quantity is measured over a process's sequences which produces a quantity of the process's internal states. The entropy rate is a useful quantity because it gives the appropriate measure of entropy of a system, its inherent randomness, given its ϵ-machine.

Statistical complexity is what Crutchfield calls his measure of structure. Understood intuitively, the statistical complexity of a process is the amount of a process's past that is determinable from the process's present state, or how much information is stored in the system. This is analytically equivalent to the amount of information that can be predicted about future states from the present state. The coin flip process has no structure in this sense because each present state of the coin allows one to determine zero information about the process's previous or future states. On the other hand, consider the so-called Golden Mean Process, where the symbol alphabet consists of 0s and 1s and random transitions from one to the other, with the exception that there can not be two 0s in a row. If the last output was a 0, one knows that both the previous and subsequent output has to be a 1, which is to say the process has structure because it contains information about past and future states within its present state.

The ϵ-machine representation is a kind of Hidden Markov Model (HMM) because the causal states of a system are themselves not directly observable, but only inferable from

[4]For further introduction to Crutchfield's computational mechanics framework, see [12, 13, 14]. However, the introductory content presented in this paper is intended to be more accessible to those who are not familiar with information theory.

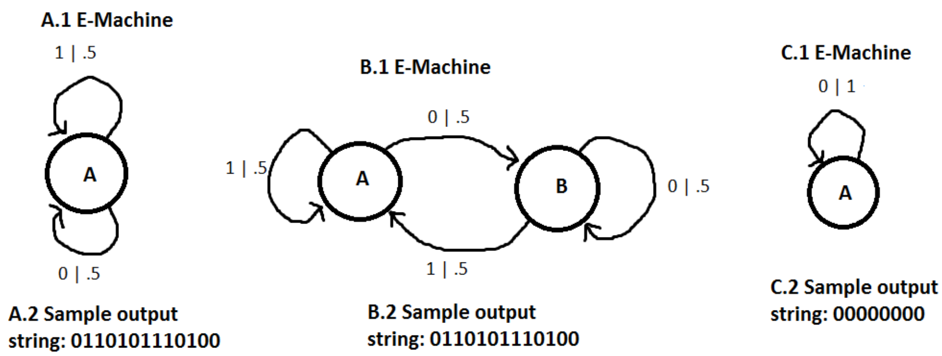

Figure 1: Pictured are ϵ-machine representations of systems and examples of the kinds of strings these systems output. The circles represent causal states and the arrows represent the transition probabilities between the causal states. The causal states are labeled alphabetically solely to help the reader visually navigate their organization. The transitions are labeled with two numbers: the first is the output symbol and the second is the probability of that transition occurring when the system is in that causal state. For example, in A.1, at any moment the machine has a 0.5 chance of outputting a 1 or a 0 in its next time-step. The numerical outputs of ϵ-machines are just convenient symbols and should not be taken to have any other mathematical qualities. **A**: A representation of an ϵ-machine for a fair coin. Notice that the system is always in the same state because at all times the system has the same relationship to future and past states, which is a total lack of structure or correlation. **B**: This ϵ-machine is structurally identical to A.1 because it gives the same output distribution, but it is represented in two states instead of one. Because B.1 can be reduced to A.1, A.1 is the the more concise, and in this case the minimal, representation of the fair coin process. In principle, any process can be represented faithfully by an infinite number of different ϵ-machines, but the minimal representation is usually visually preferable and is in fact required for proper calculations of the properties of these systems. **C**: An ϵ-machine for a string of all 0s, which could be imagined as a coin with tails on both sides.

the output values [13]. Such HMM's more closely model the relationship between observer, measurement, and system, as oftentimes the causal state of a system, physical or digital, being studied can only be inferred given measurements of the outputs of the system.

The question of how much structure a system has is separate from the question of how easily one can ascertain that structure. For example, imagine looking at a bit-string one bit at a time and trying to build an ϵ-machine to accurately model the process that created the bit string. If one came across the sequence 10101110100, after the first four bits one might think the ϵ-machine was a period-2 process, like B.1 in Figure 2. After the first ten bits, the Golden Mean Process may seem like the right representation, but the double-zero at the end invalidates that hypothesis. After a long enough sequence, perhaps it becomes clear that the system is totally random, as would be generated by a fair coin being flipped.

In such a way an observer can construct an ϵ-machine of a system from that system's output bit string, or in a more general sense can construct a representation of a system from observation. There are various algorithms that can methodically build ϵ-machines from bit strings [15, 16, 17, 18].

When I say that for a coin flip-generated bit string each newly observed symbol provides no information about the previous or subsequent symbols, *I am speaking about the structure of the system as if the ϵ-machine of the system was already known to us as observers.* This is how systems and their ϵ-machines are discussed throughout this paper. This may seem like an unusual way to talk about structure in systems, but what this does is simplifies and clarifies the questions of structure that are of concern here—namely how much structure a system has and why—by assuming the representation is already known and setting aside here the additional question of how representations are built from observations.

Given a particular ϵ-machine of a system, the statistical complexity of the system can be calculated using the formula:

$$C_\mu = -\sum_{\sigma \in S} Pr(\sigma) log_b Pr(\sigma)$$

Where C_μ is the statistical complexity and $Pr(\sigma)$ is, like above, the probability that the system is in causal state σ at any one time. The equation closely resembles the information entropy equation of a system, but when this is calculated over the ϵ-machine representation it gives a measure of the structure of a process instead of a measure of its entropy.

Given a particular number of states of a process, the most structure that process can have is one where it visits each state exactly once in a sequence before repeating, as in ϵ-machines B.1, C.1, and F.1 in Figure 2. If one knows the present state of such a process, one knows totally all the previous and subsequent states of the repeating sequence. For the same number of states, the amount of structure is minimized when every state is linked with equal probability to each other state (as in a coin or die) because there is zero information that can be determined about the state's past or future states from the present state.

Worth paying special attention to is the structure in the E.1 ϵ-machine in Figure 2 as compared to the C.1 ϵ-machine. Intuitively, E.1 appears more structured and complex than C.1. The E.1 output string is somewhat interesting in that there are discernible patterns amid otherwise random elements. However, it should be remembered that ϵ-machines are representations of systems used to analyze the structure and randomness of systems. The ϵ-machines themselves are not the systems. With structure being here defined as stored

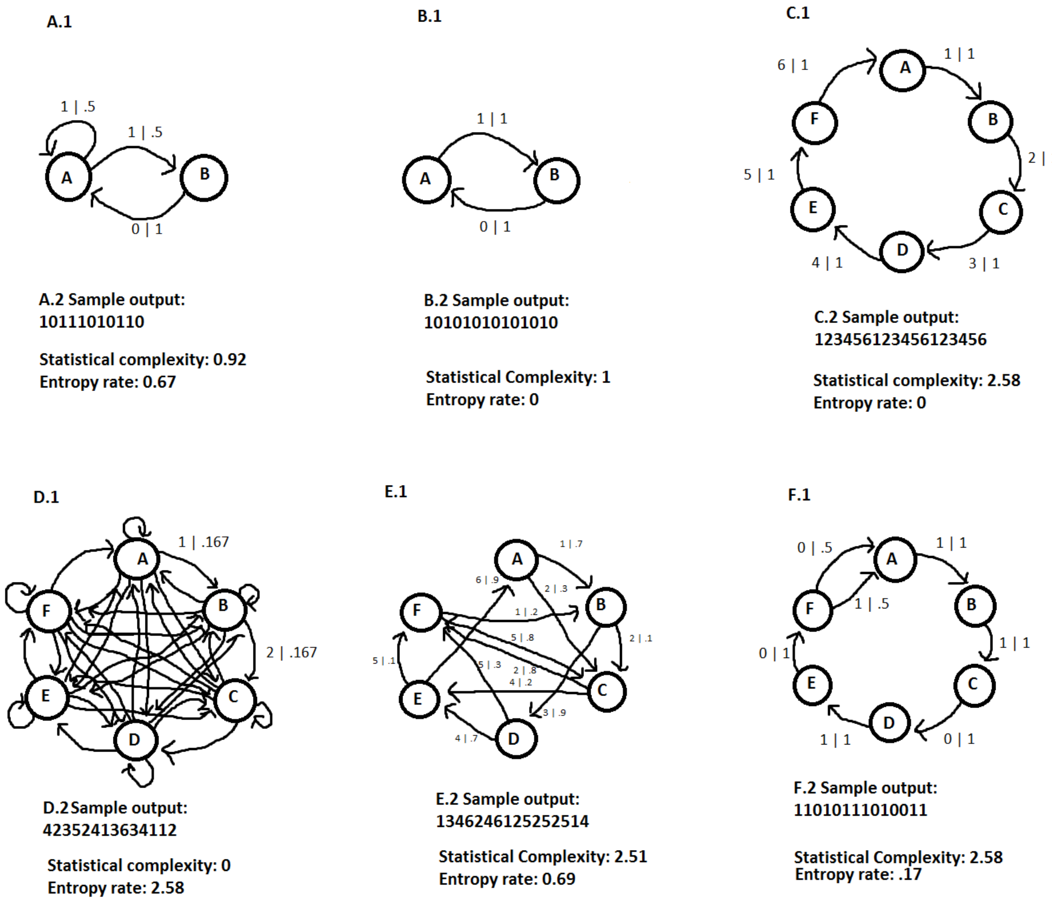

Figure 2: The aim here is to impart an intuitive feel for how ϵ-machines relate to structure and randomness. The quantities above are given in bits. **A**: The Golden Mean Process. **B**: A period-2 process. **C**: A period-6 process. **D**: The ϵ-machine for a six-sided die, where after each roll, the subsequent roll is equally likely to output any of the six possible symbols. This is the most entropic and least structured possible 6-state ϵ-machine. While this is an accurate ϵ-machine for a die, it is not the minimal one. Such a minimal ϵ-machine would consist of one state with six equiprobable transition probabilities that always leave and return to the same state, similar to coin-flip ϵ-machine in Figure 1.A.1 but with six transitions instead of two. **E**: An ϵ-machine with the deceptive appearance of having more structure than C.1. **F**: In the previous examples, the number of states are matched with the number of symbols for ease of visualization. Here there are 6 states and only 2 symbols, illustrating that it's not the number of symbols that determines the amount of structure but purely how much information is retained by the system. This is the minimum representation because the ϵ-machine can not be reduced to having fewer states.

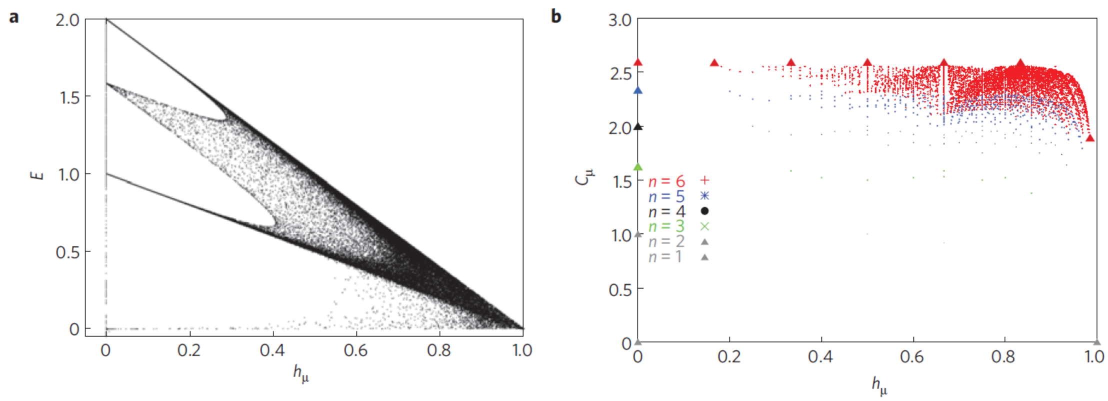

Figure 3 | Complexity-entropy diagrams. a, The one-dimensional, spin-1/2 antiferromagnetic Ising model with nearest- and next-nearest-neighbour interactions. Reproduced with permission from ref. 61, © 2008 AIP. **b**, Complexity-entropy pairs (h_μ, C_μ) for all topological binary-alphabet ϵ-machines with $n = 1, \ldots, 6$ states. For details, see refs 61 and 63.

Figure 3: Taken from [14].

information, E.1 actually stores much less information than C.1 because at any point in a binary string produced by C.1, one knows with total certainty the previous and subsequent symbols of the string given a present symbol, which means that maximum information is 'stored' in the process given its number of states. In contrast, if a digit of a string produced by E.1 is observed, one can make statistical guesses about previous and subsequent causal states, but the stochasticity of the transitions between states washes out the stored information the further from the present state one tries to ascertain its past or future. The entropy in the inter-state transitions leads to a loss of information about what can be ascertained about future and past states from the present state. Part of what makes E.1 look more complex is the way it would require inference and ingenuity to reconstruct the ϵ-machine merely from a bit string, while reconstructing C.1 from a bit string is trivially simple. As was noted above, the question of how difficult it is to ascertain a system's structure is different from the question of how much structure it has. Another part of what makes E.1 appear complex is the varying degrees of randomness within the transition probabilities. But here too, that is merely an indication of the system's entropy rate, not its structure.

Crutchfield enumerated all possible two-symbol ϵ-machines of up to six states and plotted them according to their statistical complexity and entropy rate, clearly showing how these two quantities of a process vary largely independent of each other (see Figure 3b). The exception to this independence is when the entropy rate is at or near its maximized value of 1, where each state has even or nearly even transition probabilities to itself and all other states. At a maximum level of entropy, all possibility for structure in the system washes out. This is seen in Figure 3b where red dots drop off at the right side of the complexity-entropy diagram.

While the formalities of the definition of structure above may appear removed from the intuitive sense we have of the concept of structure, when looked at closely it aligns well with our intuitions while providing a mathematical explanation for why it does so [5]. At least for the systems to which this framework can be directly applied, Crutchfield's work presents a strong case that the computational mechanics definition of structure is one built upon legitimate first principles and one for which our intuition has always been an imperfect compass.

2.3 Challenges

Before applying this framework to real world systems, it's worth discussing the issues of dimensionality and scalability.

With structure well-defined in 1-dimensional systems of bit strings, one of the primary challenges for this framework is theorizing how to apply this metric to 3-dimensional real-world systems. Crutchfield has developed an approach of laying 1D structure-quantifying methods over 2D system representations [19]. This results in structure measurements that track our intuitive ideas of structure for systems like 2-dimensional ferromagnetic Ising models and could in theory be applied to any number of dimensions, though it's applicability is still confined to systems that are grid-like in their simplicity.

There have been some notable applications of computational mechanics to real-world systems that can be made tractable by representing their output as 1-dimensional, such as neuron spiking trains [20], cellular automata [21, 22, 23, 24], macrovariables in statistical mechanical systems [23], geomagnetic data [25], multi-level time-scales in conformational dynamics of protein folding [26], fluorescence resonance energy transfer in magnesium [27], and stochastic crystals [28, 29, 30]. Working out the pragmatics of applying this framework more completely to messier 3-dimensional systems is still a long-term theoretical and experimental project. However, acknowledging these technical difficulties doesn't itself invalidate the ideas of computational mechanics or prevent them from being reasoned about and carefully applied to such messy real-world systems as far as present theory allows.

Another challenge for this theory is how to apply it to different scales and levels of description. As Shalizi and Moore show [31], computational mechanics is actually well-adapted to meet this challenge. For all measurably different states of a system, one can clump them together into larger or smaller groups of states which may act as single causal state nodes within an ϵ-machine representation, constituting coarser or finer amounts of detail. The fewer and coarser the states of an ϵ-machine for a given system, the less accurate the model becomes, but the more numerous and finer one makes the states of the ϵ-machine, the more computational power and measuring precision is required. Thus, the framework scales up and down fairly well given typical trade-offs of resources and accuracy.

3 Why Theorize Structure?

While Crutchfield's work has inspired and enabled the ideas developed below, he sticks mostly to the hard science of structure. I am interested in more general implications of Crutchfield's work, namely those stemming from what possible advances lie before us if structure truly is a fundamental property of physical systems. In this section I argue that there are reasons to believe that structure is an important explanatory property of systems. This serves to justify engaging in more conceptual and less mathematical investigations of the role of structure in systems, of which the primary hypothesis below is an example.

3.1 Structure as a Unifying Concept in the Study of Systems

A physical system consists of its material parts and the arrangement of those parts. In different parts of physics–materials science, cosmology, quantum mechanics, complexity science, etc...–the emphasis is often on finding a mechanism that explains an observed system behavior. As has been noted [32], there's no pre-existing universal method of categorizing the structure of systems in physics or one that grounds itself in first principles.

The approach to structure created by Crutchfield and Co. has the potential to be a unifying explanatory theory in the study of physical systems because it identifies structure as a property that can be measured and categorized in all systems.

3.2 Laws of the Behavior of Structure

Structure in Crutchfield's framework is a property characterized by how a system changes over time according to the causal states it transitions between. A further level of inquiry using this framework would ask how and why structure itself changes in systems over time. Just as structure is a property that tracks the degree to which systems change in non-random ways over time, so can structure itself change over time in a system in non-random ways. The following argument suggests that structure is a system property whose own behavior is governed by universal laws.

- Proposition #1: Relational properties are properties of systems that are defined by the relations between parts of a system and not by the properties of the parts themselves or by the mere summation of the properties of those parts.
  - For example, non-relational properties include mass, volume, and temperature.
- Proposition #2: Fundamental relational properties are those relational properties that play essential roles explaining the behavior of systems, and without which we don't know why a system does what it does.
  - What makes the property fundamental is that the property can not be broken down further while maintaining its explanatory power, that the property is not a mere aggregate of other properties, and that the property is not just specific to some kinds of systems but not others.
- Proposition #3: Free energy, energy, and entropy are properties that are relational and are essential to explaining the behavior of systems.
  - These concepts are relational because taking a system with the same parts, one can arrange the parts differently in such a way that changes the system's quantities of free energy, energy, and entropy.
  - These concepts are essential because without them, we have almost no general explanation for the behavior of these systems and instead all attempts at explanation must be made using particular aspects of the system instead of general concepts. Without general concepts, explanation has very little traction beyond the narrow scope of the system or set of systems being studied.

    - Regarding the essential-ness of properties, for example, chemical energy and kinetic energy are not fundamental properties but energy is. Likewise, heat is a form of entropy but not the only form.

- Conclusion #1: Free energy, energy, and entropy are fundamental relational properties.

- Proposition #4: Structure, of the kind defined by Crutchfield and Co., is relational and plays a foundational role in explaining the behavior of systems.

    - What's being posed here is not that we haven't understood system behavior because we haven't until recently had a universal and quantifiable concept of structure. Rather, what has stood in for these explanations has always been particular kinds of structured properties of systems relating particular kinds of systems. What is being posed here is that all structure belongs to a general concept of structure which in general is central to our understanding of system behavior.

- Conclusion #2: Structure is a fundamental relational property.

    - It may appear bold to elevate structure to the same conceptual level as energy and entropy in terms of how fundamental they are to understanding system behavior, but once it is clarified that it is all structure and structure in general that is being proposed, it becomes difficult to see how this could be otherwise.

- Proposition #5: Fundamental relational properties relate to each other in systems according to universal principles.

    - Anecdotal evidence that structure relates to other fundamental relational properties according to universal principles includes the following: There exists universal principles underlying relations between energy, entropy, and free energy in systems, which are themselves fundamental relational properties. The laws of thermodynamics are some such universal principles. Another one is the Gibbs Free Energy equation: free energy = the enthalpy - the temperature * the entropy. If structure is also a fundamental relational principle, then it would be in kind for it to interrelate with these other properties according to universal principles.

- Conclusion #6: Therefore the changes in structure in a system are governed by universal principles.

    - The primary thesis of this paper developed below proposes a universal principle governing the relationship of structure to free energy in systems.

The potential problems with this line of argument are many. First, Crutchfield's theory of structure may be either an incorrect theory of structure, structure itself may just not be a fundamental property of systems, or my interpretation of this property may be mistaken. Second, there's no logical necessity that fundamental properties must relate to each other in ways governed by universal principles. Or perhaps structure and what I take to be other fundamental properties exist in separate domains that don't interact after all. Beyond the possible faulty assumptions, it's possible that even if this line of reasoning is sound, actually studying how structure changes in systems over time may be too advanced for existing theory and empirical tools to make any sense of.

These difficulties notwithstanding, the potential theoretical rewards that are hinted at by such an approach to structure, the lack of any conclusive or obvious fatal counter-arguments against this approach, and the tentative conclusions reached below using this approach provide encouragement for pursuing this theoretical investigation.

3.3 Heuristic concepts of structure

The core concepts of computational mechanics are robust enough to permit their use as heuristics for reasoning about relative amounts of structure where precise quantification is not currently possible. There are innumerable systems that resist simple representation as ϵ-machines. For example, convecting liquid has patterns visible to the naked eye that impart a conviction that periodic structure is present, but modeling even a simple convection system as an ϵ-machine is not easy because its behavior does not translate straight-forwardly into a linear sequence of symbols. Also, it's unclear how to identify discrete states of a convecting liquid to act as nodes in such an ϵ-machine.

However, the general principles of computational mechanics can be applied using heuristics even when systems don't fit the mold of precise ϵ-machine representation. Formulating structure in this way expands the range of systems that computational mechanics can be used to reason about. For my purposes, it also makes tractable questions about how structure behaves in systems in general and how structure relates to other thermodynamic quantities.

These generalized tools of analyzing structure can be articulated in the same interpretive language used by Crutchfield: the more information a system stores and the more informative relations there are between parts of a system, the more structure it has. Within the computational mechanics framework, a system contains structure if one can 1) know the system's ϵ-machine representation, 2) observe some output symbols, and therefor 3) make informed predictions about future and past output symbols. Generalizing that schema beyond cases where precise ϵ-machine modeling is possible via detecting discrete states of the system and being able to represent its behavior as a linear string of symbols, we can create a rough heuristic for identifying structure in a system: a system contains structure if one can 1) roughly know the structure of the whole through its possible states and transition probabilities, 2) observe the present state of the system, and therefore 3) acquire information about future or past states of the system. This can then be used as a heuristic to ascertain whether a system contains structure because it reveals whether a system has informative relations between its parts.

4 A Specific Case of Structure Formation: From Conduction to Convection

The primary hypothesis contained in this paper concerns the process whereby structure is formed. Before postulating a general theory for this phenomena, it's worth first considering

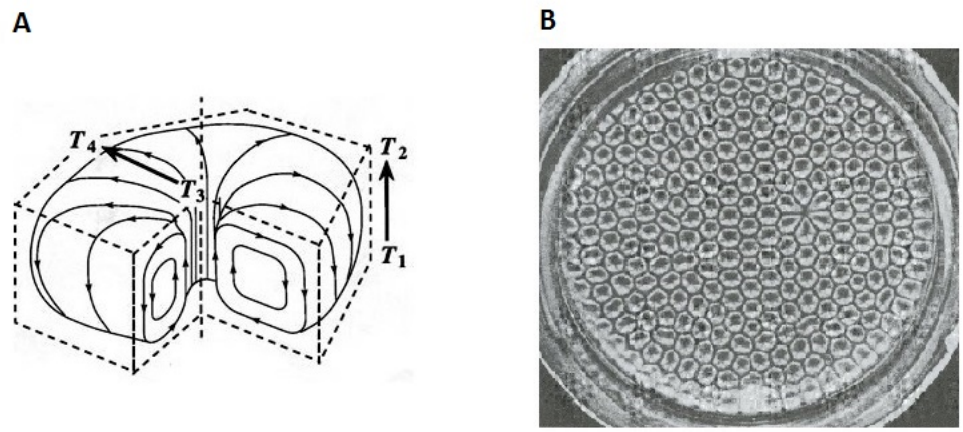

Figure 4: **A** shows the direction of the convection flows of heated liquid in a single Benard cell. **B** is a photo of Benard cells taken from above.

a well-known specific case of structure-formation, that of the transition from conduction to convection in a heated liquid.

Imagine a liquid on a hotplate being heated from below, with open air above acting as a heat sink. When the temperature gradient is small enough, heat dissipates from the plate through the liquid to the air via conduction alone. When the temperature gradient is increased by heating the plate above a certain point, the liquid begins to transport heat by convection because the hotter liquid closer to the heated plate expands and becomes more buoyant. A tension arises when the hotter liquid below is more buoyant than the cooler liquid above, and this tension is resolved when convection currents emerge that shuttle the more buoyant hotter liquid to the surface, where heat is dissipated more quickly into the air, and then the cooled liquid becomes denser and sinks downward where it is heated again and so on. In this paper, the well-known Benard cell convection will be used as the model system (see Figure 4), though all convection operates by the same general principles.

When a liquid dissipates heat by conduction alone, there is no structure being formed in the system in response to external heating. At the macro-level, there are no movements of the parts of the system, and at the micro-level of atoms, there is movement of the parts but in a totally random way. However, when a liquid dissipates heat by convection, structures are formed in the system as the parts take on specific relations to each other. Convection as a form of heat transport always occurs simultaneously with conduction, but only the former constitutes a structured process.

One can construct a rough ϵ-machine of convection to illustrate that it has structure within the framework of computational mechanics. Imagine one of the convection cycles being quartered in a cross-section, as shown in Figure 5, and that each quarter represents a causal state of an ϵ-machine. If one isolated a particle within the liquid and recorded the position of the particle at certain time-intervals, one would expect the output string of the system to be 123412341234... Given that a real-world system has stochasticity in it, the repeating pattern may occasionally be broken, with a 1 transitioning to a 3 for example, but degrees of randomness can easily be incorporated into the ϵ-machine representation. If the

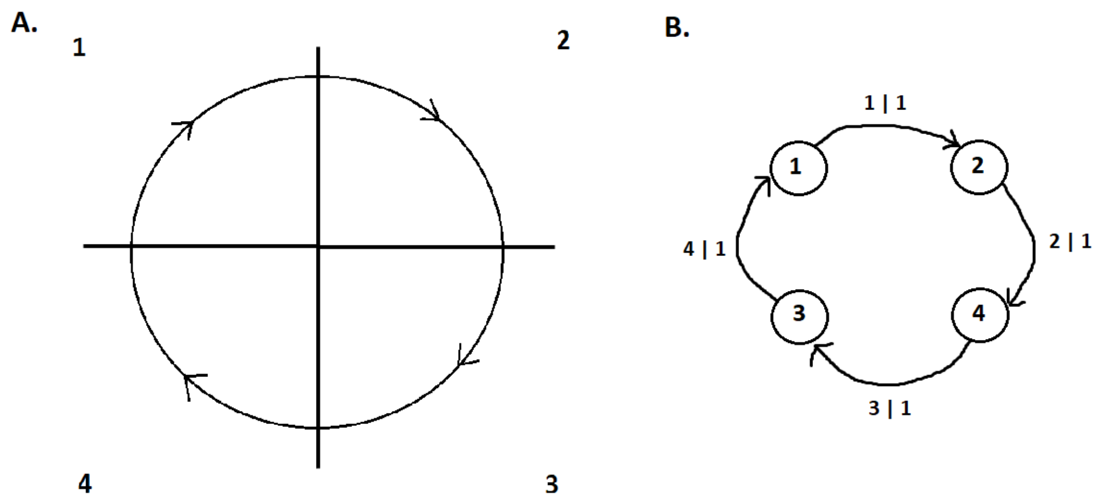

Figure 5: **A**: A representation of a path taken by a particle within a convection cycle. The cycle is quartered to track movement of the particle. When the quarters are defined as causal states of an ϵ-machine, the periodic output '1234' is indicative of structure in the system. **B**: The quadrants of the convection cycle modeled as causal states in a ϵ-machine.

same volume of water were analytically quartered and instead subjected to a heat source that merely caused conduction, the location of any particle would be almost entirely random from one time step to the next, with the output string of an ϵ-machine always in the same causal state with 4 equiprobable transitions between states, like a 4-sided die.

One could try to quantify exactly how much structure a fluid convection cycle has, but given the paucity of theory to support and guide such an attempt, that would be premature. The heuristic tools developed here from Crutchfield's framework are only powerful enough to make basic claims about whether a system has structure or not, or whether a system has more or less structure than a comparable system.

Looking at convecting fluid is a simple way to illustrate the presence of structure in a simple system, and it accurately demonstrates how the system stores information by showing that the particle cycles through the various quadrants. Recall that the ϵ-machine is not the structure itself, but just a representation that mathematically measures the structure of a system. The particle itself does not store information, the system as a whole does, and representing the system in this way illuminates how the system is structured.

4.1 Diffusion vs. dissipation

The shift from conduction to convection can be generalized in the following way. Under conduction, energy is dissipating evenly in all directions, other factors being equal. Under convection, energy is being dissipated unevenly but in a way that increases the overall rate of dissipation. Because convection has upcurrents of hot liquid and downcurrents of cooler liquid, there is an uneven, or non-uniform, spatial distribution of the way heat energy is being dissipated outward from the heat source.

Any spatially even dissipation of energy across gradients I'll call *diffusion*, whereas I'll use the term *dissipation* more generally to include both even (as in conduction) and uneven (as in convection) dissipation of energy. Diffusion of energy can occur not just via heat, but with any form of energy[5]. For example, electromagnetic energy is diffused evenly in all directions by a light bulb. Kinetic energy is likewise diffused, roughly, when a bomb explodes.

The change from diffusion to uneven dissipation occurs in a system when the energy gradient is pushed above a certain threshold by increasing the rate of energy input into the system. This threshold is passed when the parts of the system reach a critical point where diffusion breaks down because the properties of the system above that point change in such a way that the system's spatially even dissipation is no longer possible. When the diffusion breaks down, small instabilities form that gradually grow into larger instabilities that then become structures [33]. In the conduction-convection example, the increase in buoyancy of the heated liquid on the bottom is the particular pressure that breaks the diffusion process, which then produces instabilities as hotter liquid on the bottom starts to rise in columns alternating with columns of sinking colder liquid. The ultimate result is uneven dissipation of thermal energy via convection.

I'll refer to a system that dissipates energy above the threshold at which diffusion is broken as *above threshold* for short, and conversely, I'll refer to a system that dissipates energy via diffusion as *below threshold*.

5 The General Case of Structure-Formation

If structure is a fundamental property of physical systems and is as important as the concepts of energy and entropy in understanding system behavior, then there may be simple relations between structure and other fundamental properties that determine system behavior. On an intuitive level, structure is a property of a process that is arranged in some particular way, and whenever any system does anything, doing it in a particular way is different than that process occurring in a random or uniform way. Energy diffusing outward from a source results in a pure increase in entropy for the system as a whole without any resulting increase in structure. However, when a source dissipates energy through a process that is not even, some kind of explanation is required for why the dissipation is uneven because spatially even dissipation, all things being equal, is the default behavior expected from the 2nd law of thermodynamics that demands equalizing energy gradients over time. Of course, specific explanations can be given for why energy dissipates unevenly in a particular system, but the uneven dissipation of energy seems to be a general phenomena and thus requires a general explanation.

I believe this explanation requires relating structure and energy to each other. Just as the first and second laws of thermodynamics are very simple relations regarding entropy and energy, it seems plausible that the most basic relations between structure and other

[5]Entropy production is the same concept as energy dissipation, just posed in different terms. The choice of terminology in a specific context comes down to what is more conceptually convenient or clear. In other work, I make use of recent developments in the study of entropy production, and the work here should also be seen as connected to that.

properties also are simple relations.

From both common intuition and the framework advanced here, it is clear and unsurprising that when a heated liquid system crosses the threshold from conduction to convection, structure emerges. The main thesis of this paper generalizes this effect and claims it applies to all structures and systems: *The formation of structure in systems is always and only the consequence of the dissipation of energy at a higher rate than can be accommodated by diffusion.* This is posed as a universal principle of the relation between energy and structure, which offers more evidence for proposition #5 above.

What makes the case of the transition from conduction to convection conceptually clear is that the change from the unstructured to the structured state is prompted by a single change in a single system parameter: the increase in heat applied to the system above a specific threshold. Many other systems that appear to contain structure are a less convenient testing ground for studying structure as a fundamental property because there are so many other confounding aspects of the system that make it difficult to pare down the parts of the system to the bare essentials responsible for the dynamics of structure. The conduction-to-convection example has no need of ongoing material inputs; the appearance of structure emergences in a homogeneous material; it's the kind of system that can be set up, observed, and studied using simple experimental procedures. Other prototypical structures in nature, including biological, geological, astronomical, chemical, atomic, quantum mechanical systems, are typically not so convenient.

As put forth above, diffusion is structureless, meaning that a system that is merely diffusing energy contains no informative correlations between its parts. Any break from diffusion necessarily results in a process that's structured, because any substantive change in a system's state transition probabilities not along the axis from less to more entropy necessarily pops out along the axis from less to more structure. I'll call uneven dissipation above threshold *structured dissipation*[6].

Insofar as ϵ-machines can represent the basic behavior of systems as transitioning between states, the argument that non-entropic changes in a system imply structural changes can also be posed in terms of ϵ-machines. Imagine the ϵ-machine for the coin-flip process. When the entropy of the process is altered by changing the bias of the coin and thus altering the output probabilities, no structure is added or changed. If one changes the process by adding new output symbols, for example by changing the process to the roll of a die with six possible outcomes instead of the coin flip's two outcomes, again entropy is added to the system but the structure remains unchanged.

There are only two possible ways to alter an ϵ-machine without altering the ϵ-machine's entropy. The first and trivial way is to change the ϵ-machine in such a way that it might be symbolically or otherwise altered but still retain the fundamental properties of the ϵ-machine intact. For example, instead of a coin with 0s and 1s, it could be a coin with 1s and 2s, or a die with three 1s and three 0s, etc... Such trivial changes don't affect the entropy or the structure of the system. The second, and only substantive, way to change an ϵ-machine without altering its entropy requires making changes to its structure. Or stated conversely,

[6]Structured dissipation bears an obvious resemblance to Ilya Prigogine's theory of dissipative structures", though they are very different theories that happen to engage some similar phenomena. Furthermore, many have questioned the substance of Prigogine's contributions on this subject [34, 35, 36, 37].

any substantive change in an ϵ-machine that can't be entirely accounted for as a change in entropy necessarily implies a change in structure.

This idea can be illustrated using specific ϵ-machine examples. Taking the simplest case of an ϵ-machine that contains no structure, its minimal representation must contain only one state, as in the coin-flip ϵ-machine. If there are no informative relations between future and past states, then the ϵ-machine can give a varying range of outputs, as in the different sides of a coin or sides of a die, but those outputs can never have informative correlations between them, and therefore the system is always in the same causal state in relation to future and past states. To change the system in any way that doesn't merely adjust the output probabilities and entropy values requires adding another state to the ϵ-machine. For any minimal ϵ-machine representation, an increase in the number of causal states results in an increase in structure of the ϵ-machine, other things being equal. Such multi-state, or structure-containing, ϵ-machines must have outputs with informative relations between them because it requires that the system store information about past and future states within its present state.

Any open and stable driven system engaged only in diffusion is essentially a one-state ϵ-machine. A liquid being heated by an external source which then conducts that heat to a heat sink is always in the same macrostate. However, once convection emerges in a liquid, a description of the system needs to include the changing states of the system as the particles move through their cycles.

Notice that there is no third property of ϵ-machines whereby a system can be substantively altered in a way not relating to its structure or entropy. To substantively alter a system is to change the way it transitions between states, and the nature of these transitions is entirely captured by a description of the system's structure and entropy, as is represented in ϵ-machine diagrams. This is also a conceptual argument in support of the above proposition #5 that structure relates to other fundamental relational properties, in this case entropy, according to universal principles.

Any system capable of forming structure has some point at which it crosses the threshold from diffusion to structured dissipation. Where that critical threshold lies in a particular system is determined by the thermal properties of its materials as well as many other specific properties. What is universal to all structures is their emergence in systems driven above threshold.

6 Studying Structured Dissipation in Other Systems

The transition from conduction to convection in a liquid appears like a straight-forward case of patterns and structure emerging from a system that previously lacked pattern or structure. Many supporting arguments and assumptions need to be validated for this hypothesis to hold up but various points make it plausible and worth testing elsewhere.

How might the hypothesis of structured dissipation be explored in other systems? Finding systems that go from an unstructured state to a structured one in response to an increase in energy input is not difficult. One example is the fusion of atoms into heavier elements in stars once those stars reach a certain mass with enough gravitational energy to create extreme internal pressures that stimulate nucleosynthesis. Another is the kind of chemical cycles that synthesize more complex molecules in response to sufficient energy flux. Turbulence is another prime example when discrete microstructures in the form of vortexes are observed upon the transition from laminar flow to turbulent flow in response to increased flow rate, something which Crutchfield and colleagues have been working on applying their framework to already [38, 39].

The difficult part is then fitting these examples of the emergence of structure in such a way that one can show that the structure observed can be justified as being structure in the mathematical meaning used in computational mechanics. Because the mathematical framework is not developed far enough yet to easily accommodate such systems, one has to use other methods to reason about the change in structure in a system over time. In this paper, I use some heuristics based on computational mechanics concepts.

After sample systems are found to fit the criteria necessary to apply this specific definition of structure, the theory of structured dissipation can be tested. For example, does structure in systems only emerge after an increase in energy input above a certain threshold and below which energy diffuses evenly? Does the emergence of structure facilitate a faster diffusion of energy overall? Does structure ever emerge in these systems without a concomitant increase in energy dissipation? These and other questions can help make the theory of structured dissipation falsifiable.

There are some common traps to avoid in using this approach. Often what appears to be the emergence of structure in a system is found to be merely the emergence of order. Like the sorting of different sized grains or pebbles in a container such that the smaller ones go to the bottom and the larger ones go to the top, this appears at first glance to be a structured process, but upon closer inquiry this constitutes mere ordered sorting and does not qualify as structured in the sense of computational mechanics. The formation of many kinds of apparently complex and structured crystals and other materials also exhibit this kind of ordering.

Structure in the colloquial sense can be applied to both objects and processes. However, computational mechanics itself can likely only be applied to structured processes and not structured objects [23]. There is no coherent way of taking a static object that appears to exhibit structured patterns and trying to fit it into an ϵ-machine diagram because a static object in no way transitions between different states as required by ϵ-machine modeling. However, computational mechanics would be a pretty weak theory if it was not able at all to account for the kind of structure observed in objects, as such structures–from atoms and molecules to geological formations–constitute a major portion of our intuition of what structure is. The key to analyzing object structures through this lens is to recognize that for an object structure to have been created, there must have been a structured process that created it, and so one can go about applying computational mechanics concepts to the structured processes that create structured objects.

7 Theoretical and Practical Uses for a Concept of Structure

In the history of physics, whenever a new fundamental concept like energy or entropy has been formalized and quantified, it has opened up new domains of inquiry and been accompanied by far-reaching and influential theorizing. The formalization and quantification of structure could have a similar effect on science. Just like energy and entropy are used universally across the hard and life sciences, so too could a concept structure find broad application.

One possible application is to the study of the emergence of life in the lab and on other planets, which is in many ways a study of the emergence of a broad class of structures including phenomena like biomolecules, proto-cells, and simple organisms. Applying a more formal concept of structure to such studies may make clear new constraints and reveal new possibilities about how life emerges from non-living matter. The study of many various kinds of complex phenomena, from brain activity to stock markets, could perhaps benefit from new modeling tools and conceptual frameworks like those developed in computational mechanics. However, speculating much on the specific details of such possibilities in science in general is beyond the reach of my knowledge and imagination.

For industrial applications, just as the concepts of entropy and energy are used universally by engineers to maximize efficiency, so too could structure perhaps play a similar role in some production processes. Of course, the structure of a hammer or car is not what structure in the sense of computational mechanics is useful for. But in terms of certain complex manufactured materials possessing structures, finding ways to maximize the efficient production of structure may result in new cost-effective techniques.

8 Conclusion

Recent advances in the definition of physical structure enable us to begin exploring the science of structure in real world systems. Using Crutchfield's computational mechanics framework as a basis for understanding structure, this paper has argued that structure is always and only the consequence of a system being driven above the threshold at which it can diffuse energy evenly in all directions. I call this structured dissipation, and I propose it to be a universal principle of structure in physical systems.

Even when precise quantification of structure in systems is impossible, the ideas of computational mechanics permit the use of heuristic concepts of structure to reason about comparative amounts of structure in a system under different conditions and in different states. With this form of reasoning, structured dissipation can be tested by looking at specific systems. The emergence of convection cycles in a heated liquid is taken to be a model case for illustrating this principle.

The ideas presented here are exploratory and require much work, particularly on the gap between clean quantitative models and messy real-world systems. However, a mathematical formulation of structured dissipation seems entirely plausible. If such progress were forthcoming and other basic principles of the behavior of structure were discovered, it is exciting to think about the scientific potential of a science of structure dynamics for advancing our understanding of physical systems.

References

[1] J. P. Crutchfield. "The Origins of Computational Mechanics: A Brief Intellectual History and Several Clarifications." *arXiv preprint arXiv:1710.06832*, 2017

[2] P. Caws. *Structuralism: A Philosophy for the human sciences*. Humanities Press, 1997.

[3] B. Russell. *Introduction to mathematical philosophy*. Macmillan, 1920.

[4] C. E. Shannon. "A Mathematical Theory of Communication." *Bell System Technical Journal*, 27: 379-423, 623-656, 1948.

[5] J. Ladyman, J. Lambert, K. Wiesner. "What is a complex system?." *European Journal for Philosophy of Science*, 3(1): 33-67, 2013.

[6] E. T. Jaynes. "Information Theory and Statistical Mechanics." *Physical review*, 106(4), 620, 1957.

[7] J. M. Parrondo, J. M. Horowitz, T. Sagawa. "Thermodynamics of Information." *Nature Physics*, 11(2): 131-139, 2015.

[8] D. P. Feldman, C. S. McTague, J. P. Crutchfield. "The Organization of Intrinsic Computation: Complexity-Entropy Diagrams and the Diversity of Natural Information Processing." *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 18(4): 043106, 2008.

[9] S. Wolfram. "Statistical Mechanics of Cellular Automata." *Reviews of Modern Physics*, 55: 601-644, 1983.

[10] Y.L. Klimontovich. *Turbulent Motion and the Structure of Chaos: A New Approach to the Statistical Theory of Open Systems*. Kluwer Academic, 1991.

[11] C. R. Shalizi, J. P. Crutchfield. "Computational mechanics: Pattern and prediction, structure and simplicity." *Journal of statistical physics*, 104(3-4): 817-879, 2001.

[12] J. P. Crutchfield, K. Young. "Inferring statistical complexity." *Physical Review Letters*, 63(2): 105, 1989.

[13] J. P. Crutchfield. "The calculi of emergence: computation, dynamics and induction." *Physica D: Nonlinear Phenomena*, 75(1): 11-54, 1994.

[14] J. P. Crutchfield. "Between order and chaos." *Nature Physics*, 8(1): 17-24, 2012.

[15] J. P. Crutchfield. "Information and its metric." *Nonlinear Structures in Physical Systems*, Springer New York, 119-130, 1990.

[16] J. P. Crutchfield, D. P. Feldman. “Synchronizing to the environment: Information-theoretic constraints on agent learning.” *Advances in Complex Systems*, 4(02n03): 251-264, 2001.

[17] N. F. Travers, J. P. Crutchfield. “Exact synchronization for finite-state sources.” *Journal of Statistical Physics*, 145(5): 1181-1201, 2011.

[18] J. P. Crutchfield. “ϵ-Machine Reconstruction.” *Natural Computation and Self-Organization: The Physics of Information Processing in Complex Systems*, Lecture 23, `https://www.youtube.com/watch?v=vnUv430ROCA`, 2013.

[19] D. P. Feldman, J. P. Crutchfield. “Structural information in two-dimensional patterns: Entropy convergence and excess entropy.” *Physical Review E*, 67(5): 051104, 2003.

[20] S. E. Marzen, M. R. DeWeese, J. P. Crutchfield. “Time resolution dependence of information measures for spiking neurons: Scaling and universality.” *Frontiers in computational neuroscience*, 9, 2015.

[21] J. P. Crutchfield, J. E. Hanson. “Turbulent pattern bases for cellular automata.” *Physica D: Nonlinear Phenomena*, 69(3): 279-301, 1993.

[22] J. P. Crutchfield, J. E. Hanson. “Computational mechanics of cellular automata: An example.” *Physica D: Nonlinear Phenomena*, 103(1): 169-189, 1997.

[23] C. R. Shalizi. *Causal architecture, complexity and self-organization in the time series and cellular automata*. Doctoral dissertation, University of Wisconsin–Madison, 2001.

[24] C. R. Shalizi, R. Haslinger, J. B. Rouquier, K. L. Klinkner, C. Moore. “Automatic filters for the detection of coherent structure in spatiotemporal systems.” *Physical Review E*, 73(3): 036104, 2006.

[25] R Clarke, M Freeman, N Watkins. “The application of computational mechanics to the analysis of natural data: An example in geomagnetism.” *Phys. Rev. E*, 67: 160203, 2003.

[26] C-B Li, H Yang, T Komatsuzaki. “Multiscale complex network of protein conformational fluctuations in single-molecule time series.” *Proc. Natl Acad. Sci. USA*, 105: 536541, 2008.

[27] D Kelly, M Dillingham, A Hudson, K Wiesner. “A New Method for Inferring Hidden Markov Models from Noisy Time Sequences.” *PLoS ONE*, 7(1): e29703, 2012.

[28] D. P. Varn, G. S. Canright, J. P. Crutchfield. “Discovering planar disorder in close-packed structures from x-ray diffraction: Beyond the fault model.” *Physical Review B*, 66(17): 174110, 2002.

[29] D. P. Varn, G. S. Canright, J. P. Crutchfield. “ϵ-Machine spectral reconstruction theory: a direct method for inferring planar disorder and structure from X-ray diffraction studies.” *Acta Crystallographica Section A: Foundations of Crystallography*, 69(2): 197-206, 2013.

[30] D. P. Varn, J. P. Crutchfield. "Chaotic crystallography: how the physics of information reveals structural order in materials." *Current Opinion in Chemical Engineering*, 7: 47-56, 2015.

[31] C. R. Shalizi, C. Moore. "What is a macrostate? Subjective observations and objective dynamics." arXiv preprint cond-mat/0303625, 2003.

[32] M. C. Cross, P. C. Hohenberg. "Pattern formation outside of equilibrium." *Reviews of modern physics*, 65(3): 851, 1993.

[33] J. P. Gollub, J. S. Langer. "Pattern formation in nonequilibrium physics." *Reviews of Modern Physics*, 71(2): S396, 1999.

[34] J. Keizer, R. F. Fox. "Qualms regarding the range of validity of the Glansdorff-Prigogine criterion for stability of non-equilibrium states." *Proceedings of the National Academy of Sciences*, 71(1): 192-196, 1974.

[35] J. Horgan. "From Complexity to Perplexity." *Scientific American*, 272: 104-109, 1995

[36] P. W. Anderson, D. L. Stein. "Broken symmetry, emergent properties, dissipative structures, life." In *Self-organizing systems*, 445-457, Springer US, 1987.

[37] J. Bricmont. "Science of chaos or chaos in science?." *Annals of the New York Academy of Sciences*, 775(1): 131-175, 1995.

[38] A. Rupe, J. P. Crutchfield, K. Kashinath, Prabhat. "A Physics-Based Approach to Unsupervised Discovery of Coherent Structures in Spatiotemporal Systems." *eprint arXiv*, 1709.03184: 1-4, 2017.

[39] A. Rupe, J. P. Crutchfield. "Local causal states and discrete coherent structures." *eprint arXiv*, 1801.00515: 1-27, 2018.