

Noisy Dynamical Systems Evolve Error Correcting Codes and Modularity

Trevor McCourt,^{1,2,*} Ila R. Fiete,^{3,4} and Isaac L. Chuang^{1,5,2}

¹*Department of Electrical Engineering and Computer Science, [Massachusetts Institute of Technology](#), Cambridge, Massachusetts 02139, USA*

²*The NSF AI Institute for Artificial Intelligence and Fundamental Interactions, Cambridge, Massachusetts 02139, USA*

³*Department of Brain and Cognitive Sciences, [Massachusetts Institute of Technology](#), Cambridge, Massachusetts, USA*

⁴*McGovern Institute for Brain Research, Department of Brain and Cognitive Sciences, [Massachusetts Institute of Technology](#), Cambridge, Massachusetts 02139, USA*

⁵*Department of Physics, [Massachusetts Institute of Technology](#), Cambridge, Massachusetts 02139, USA*



(Received 16 May 2023; accepted 6 March 2026; published 3 April 2026)

Noise is a ubiquitous feature of the physical world. As a result, the first prerequisite of life is fault tolerance: maintaining integrity of state despite external bombardment. Recent experimental advances have revealed that some biological systems achieve fault tolerance by implementing mathematically intricate error correcting codes and by organizing in a modular fashion that physically separates functionally distinct subsystems. These elaborate structures represent a vanishing volume in the massive genetic configuration space. How is it possible that the primitive process of evolution, by which all biological systems evolved, achieved such unusual results? In this work, through experiments in Boolean networks, we show that the simultaneous presence of error correction and modularity in biological systems is no coincidence. Rather, it is a typical co-occurrence in noisy dynamic systems undergoing evolution. From this, we deduce the principle of error correction enhanced evolvability: systems possessing error correcting capabilities are more effectively improved by evolution than those without.

DOI: [10.1103/7kcl-jshm](#)

I. INTRODUCTION

Even under the influence of entirely deterministic natural laws, a finite agent that interacts with a large environment will experience some degree of randomness. This is because it will eventually lose record of all its previous interactions, and will only be able to reason on probabilistic terms. This is exemplified in physics by observations of Brownian motion [1], in which an observed particle in contact with a large bath of unobserved particles undergoes seemingly random, yet entirely determined, motion. Therefore it is reasonable to suggest that anywhere life emerges, it is probably doing so despite a stochastic environment, as it has on Earth.

von Neumann appreciated the fundamental nature of noise and observed that the biological information processing systems that underpin life handle it elegantly [2]. In particular, biological systems are generally not halted by a single-component failure or error. In contrast, the engineered computers of his time were generally brittle, and errors in a single component had to be identified and addressed before the rest of the system could continue operation. von Neumann thought this approach would begin to become infeasible as more and more complex systems were constructed [3]. To

resolve this and to contribute further to our understanding of the connection between computing in the presence of noise and life, he established the theory of fault-tolerant computation [4].

Fault-tolerant computing systems compose noisy primitive computational units (e.g., logic gates) into structures that compute the same primitive, robustly [4,5]. This concept is not limited to digital computation: fault tolerance has recently been demonstrated in neural networks with brainlike representations [6] and is also an active area of theoretical and experimental research in quantum computing [7–10], where it is particularly important due to the high susceptibility of real qubits to noise. Today, fault-tolerant computing is generally accomplished by computing within the context of an error correcting code. An error correcting code specifies a way to robustly store or transmit messages in the presence of noise [11]. The basic idea is to encode messages in code words, which add redundancy such that if part of the code word is lost the message can still be recovered. Fault-tolerant computing further requires that the recovery be done by the noisy system itself.

Since the time of von Neumann, many specific examples of robust behavior in biological information processing systems have been discovered. Genomes are one example of this. Single-knockout experiments in a particular yeast found that only 19% of the genes were required for the organism to continue living [12,13]. The same resilience is seen in the human brain, which can tolerate substantial death of neurons with minimal cognitive impairment [14,15]. This level of robustness implies the existence of mechanisms for error correction. Studying these mechanisms is an emerging topic

*Contact author: tmccourt@mit.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](#) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

of theoretical interest in genetics [16]. There is also substantial evidence for formal error correction mechanisms in the brain. Many neural circuits have population activities that live on low-dimensional manifolds [17–20], reminiscent of topological error correcting codes [21,22]. In particular, it is thought that grid cells use an efficient error correcting code to robustly encode position in the presence of noise [23].

One way that large-scale natural and engineered systems seem to achieve robustness is through *modularity*, the division of a large system into functionally independent parts. The genome is believed to have spatial and functional modularity [24], and a common view of the brain is that it is composed of many modules responsible for different functions that combine to form our intelligence [25]. At a large scale, organizations, processes, and machines frequently have a division of responsibility for separate subtasks, such that if one subsystem fails the system can still complete its function [26–28]. In general, an overarching motif observed in natural systems is that they appear to be assembled out of smaller, individually robust components.

At the same time, biological structures emerge by making local changes that improve fitness, a process of adaptive evolution. Robustness has often been thought to be at odds with evolution; it is intuitive that systems that are robust to perturbations will also be difficult to change and improve [29–31]. Additionally, the particular modular, error correcting structures seen in biology represent a vanishing fraction of genomic configuration space, presenting evolution with a daunting search problem. A tantalizing question is, therefore, what conditions and principles drove evolution towards these robust structures? The answer to this question is significant, as it gets to the root of forces that gave way to the crystallization of life.

Any property that increases the efficacy of adaptive evolution may be referred to as *evolvable*. Evolvable properties should emerge as a typical result of evolution since obtaining some of an evolvable property increases an organism’s ability to acquire more of it. Identifying evolvable properties is a topic of constant interest [32–35]. For example, Kauffman famously showed that systems with a phase transition from frozen to chaotic behavior (a so-called ordered phase) tend to be more evolvable, and seek the edge of chaos [36,37].

In this work, we show that both error correcting codes and modularity are typical coemergent results of evolution in a noisy environment. We show that this occurs because organisms with error correcting properties are better protected from lethal mutations than those without, allowing them to more effectively search configuration space for improvements. From this, we introduce the concept of error correction enhanced evolvability: organisms with error correcting codes are more evolvable than those without. This principle is illustrated in Fig. 1. Noise bootstraps this phenomenon, suggesting that noise plays an important role in evolving complex structures.

We begin by introducing Boolean networks and motivating their use in the study of evolution and life. We then show that when evolved to perform primitive computations in the presence of noise, Boolean networks almost always develop strong error correcting codes. By probing the effect of mutations on evolved networks, we then demonstrate that this occurs because error correcting codes increase the number

Error correction enhanced evolvability




FIG. 1. Error correction smooths valleys in fitness landscapes. Noisy systems initially evolve error correction to protect themselves from random environmental fluctuations that affect their state. However, once evolved, error correction also helps *neutralize* permanent genetic mutations that would otherwise be lethal. During evolution, this allows organisms that possess error correcting codes to search further in genome configuration space for improvements than their non-error correcting counterparts. This is illustrated in the figure: the organism that uses error correction can sustain the first genome mutation without falling into a valley of fitness that is too low to survive. It can then acquire a further mutation which brings it to a new, improved local fitness maximum, corresponding to an improved error correcting code. With this new code, it will even more easily be able to find improvements, and so on. Error correction begets more error correction, leading to its ubiquity in the natural world.

of neutral mutations seen by organisms, allowing them to more effectively search locally for improvements than their non-error correcting counterparts. Finally, we show that under composite computational tasks (that require more than one primitive), evolution tends to produce structurally modular organisms that employ modular error correcting codes, beautifully reflecting what is seen in biology.

II. BOOLEAN NETWORKS

Originally developed to model gene regulatory networks [38–40], Boolean networks are autonomous, discrete-time dynamical systems with binary-valued states. Boolean networks can be viewed as dynamic function evaluators. A function input can be given as an initial state to a network and the output is simply the state at some specified later time. Boolean networks are therefore reasonable computational models of primitive living organisms: receiving input and later producing the correct output is equivalent to an agent responding to environmental stimulus to avoid death [41]. Despite their simple binary nature, Boolean networks can express very complex behavior. For example, they possess the aforementioned ordered phase that seems to be crucial for successful adaptation [36,42]. This complexity makes them a candidate for modeling large-scale brain dynamics [43].




FIG. 2. Boolean networks adapting into error correction. (a) A Boolean network that has adapted to solve the and task in the presence of noise. Inputs are provided to the orange nodes 0 and 1 at time $t=0$ and the answer is expected at the green node 2 at $t=T$. A Boolean network is specified by a directed graph with a truth table living on each node. The state of each head node at time $t+1$ is a binary function of the state of each tail node at time t . For example, the state of node 6 at time $t+1$ is computed by applying a length 2^3 truth table to the states of nodes 0, 1, and 3 at time t . The edge coloring reflects the influence of a head node on a tail node. (b) Examples of noisy dynamical trajectories of the network in (a) for each of the 2^2 possible input states. The network encodes the answer in code words of maximal Hamming distance and stabilizes the code word against noise events. Code words may be taken as the final state of the network $x[T]$ for the different output values. The two code words are visualized on the network graph. The boundary of the circles indicates the node function, and the fill indicates the code word. (c) Evolutionary trajectories of 150 randomly initialized populations learning to solve the XOR and AND tasks. Populations were not pruned; every population that was started is included in the statistics. The solid line indicates the median logical error probability, and the shaded region indicates the interquartile range. (d) Distribution of final logical error probabilities over all populations. In both tasks, typical organisms learn to suppress errors far below the physical noise level (p_{phys} , indicated by the dashed black line). (e) Distribution of average Hamming distance between the code words C_0 and C_1 over all populations. The typical organism uses a strong error correcting code to suppress errors. (f) Scaling of the logical error probability with the physical error probability for the highest performing organisms: logical error probability is strongly suppressed relative to physical error probability.

Here, we leverage Boolean networks to computationally probe evolution in a noisy environment. We assign Boolean network organisms a measure of fitness based on their ability to compute a simple function in the presence of noise. We provide an organism with the input to the computation in a raw, unencoded form, and allow it to undergo noisy dynamics for several time steps. It is then queried for the output, which can be compared to the correct output for fitness calculation. Producing the correct output in the presence of noise is a daunting task: an organism has to have a method for encoding the inputs, robustly performing the computation, and stabilizing the answer against noise in memory. We study the evolution of organisms that compute primitive and more complicated composite computations to study the emergence of error correction and modularity.

The state of a Boolean network of N nodes is given by a length N binary vector x . The state of the i th node x_i at time $t+1$ is given as an arbitrary Boolean function f_i of the state of $k_i \leq k_{\text{max}}$ other nodes at time t . Therefore, a Boolean network can be thought of as a directed graph with truth tables of length 2^{k_i} living on each node [Fig. 2(a)]. The genome of such a Boolean network is therefore specified by a matrix

containing the truth tables for each node and a matrix defining the connectivity between nodes. For further information on how the Boolean network genome was specified in these experiments, see [44] or Supplemental Material (SM), Sec. 1 [45].

The space of possible Boolean networks is truly massive. For a network with exactly k connections per node, the number of unique networks is [46,47]

$$\left(\frac{2^k N!}{(N-k)!} \right)^N. \quad (1)$$

The networks considered in this work are *ragged*, which means that they have $k_i \leq k_{\text{max}}$ connections per node. As such, the number of networks is larger than what is indicated by Eq. (1). For the $N=8$ and $k_{\text{max}}=3$ network shown in Fig. 2(a), this evaluates to $\sim 10^{40}$. This number grows very rapidly with N ; for $N=15$ there are more configurations than there are bits of information in the observable universe [48]. Therefore, any particular configuration of a Boolean network resides in a vanishing volume of configuration space. At the same time, Boolean network loss landscapes are generally

rugged and multi peaked, making adaptation nontrivial [36]. Therefore, any configuration that is consistently reached by randomly initialized adaptive walks must be somehow attractive: it is exceptionally difficult to stumble upon particular solutions by accident when working with Boolean networks. This makes Boolean networks particularly useful for studying evolvability.

We introduce stochasticity to our Boolean networks [49] by injecting independent Bernoulli noise alongside the deterministic state update described above:

$$x[t + 1] = f(x[t]) \oplus \text{Bernoulli}(p_{\text{phys}}), \quad (2)$$

where f is the collection of all the f_i and represents the deterministic update previously described, $\oplus$ indicates bitwise XOR (exclusive or), and $p_{\text{phys}} \sim 10^{-2}$ is the physical error probability. In other words, the noise acts to flip each bit of $f(x[t])$ independently with probability p_{phys} . Figure 2(b) shows an example of the state of a noisy Boolean network versus time for different inputs.

We provide input to our organisms via designated input nodes, which are initialized with the input to a computational problem at $t = 0$. For example, an organism that solves the and task has two input nodes, as shown in Fig. 2(a), which could be initialized as any two-bit binary string. The rest of the nodes are initialized to 0. The organism then undergoes $T - 1$ steps of noisy dynamics, corresponding to deterministic state updates followed by exposure to noise [Eq. (2)]. One more state update is then applied, serving as a noiseless decoding step, as is standard in the study of fault-tolerant subsystems [4]. The state of the output nodes is then taken as the logical output of the organism. We randomize T within some small range to force the organism to stabilize the solution instead of being allowed to pass through it transiently. The structure and complexity of the computational task can be varied to probe different properties of the solutions. The unfitness (or *loss*) of a given organism is its logical error probability p_{log} . p_{log} is the probability that the output node contains the wrong value at time T (after the decoding step), averaged over noise and all possible input binary strings. As an example, for the network shown in Fig. 2(a), p_{log} is the probability that at time T the output node does not contain the and of the values initially supplied to the input nodes. p_{log} therefore reflects the effect of physical error p_{phys} on the network. For a fixed p_{phys} , a low p_{log} means the network is performing the computation in a way that is robust to noise.

Endowed with a measure of fitness, we can then experiment with adaptive evolution in configuration space.

Our evolutionary procedure follows an externally imposed “survival of the fittest” rule, as is standard in computational studies of evolution [50–54]. At the beginning of an experiment, a population of randomly generated organisms is created. Then several generations of evolution are carried out to select for organisms with increased fitness. During each generation, the fitness of each organism in the population is evaluated by measuring its logical error rate, and the least fit organisms are dropped from the experiment. The genome of the surviving organisms is used to create replacements via crossover operations (breeding) and mutation. This procedure is detailed in SM, Sec. 2 [45] and in the code we used to carry out our experiments [55].

This evolutionary procedure is intentionally primal and is meant to mirror operations that would have been feasible during biogenesis. We tried a few different crossover operators and found that none of them meaningfully outperformed the local search (see SM, Sec. 4 [45]).

Note that we are not attempting to use Boolean networks as a model of any specific biological system. Instead, we are simply using the networks as a toy example of something that can undergo an evolutionary process that is qualitatively similar to what could happen in a biological system, as has been done in many prior works [56,57].

Despite the apparent simplicity of Boolean networks, exploring their evolution in the presence of noise is extremely computationally intensive. For every evolutionary step, several hundred populations of several hundred organisms must be run through several hundred different noise trajectories starting from tens of different input states. This number quickly approaches batches of 10^8 networks. Such large experiments were made possible by accelerating evaluation using modern graphics processing units (GPUs) with large memories [44]. Trillions of networks were simulated in the course of this work, consuming several GPU months of compute time.

One reason for this computational difficulty is the discrete nature of the genome of a Boolean network. It would be reasonable to avoid this difficulty by working with a system that has friendlier optimization properties, such as a neural network, which could allow for much-larger-scale experiments. Despite the inherent limitations, there is significant value in working with a discrete system because, as far as we know, biological evolution also occurred in a discrete space.

We performed several experiments that probed the ability of Boolean networks to robustly perform the primitive computations and AND XOR in the presence of noise. Figure 2(a) shows a typical resulting organism for the and task (xor results are similar; see Supplemental Material, Sec. 3 [45]). The organism itself has little structure, as expected for a primitive task.

III. EVOLVING ERROR CORRECTION

What is most interesting about the evolved organisms is that they seem to achieve low error rates by implementing error correcting codes. There are two possible answers to both the and AND XOR tasks: 0 and 1. As shown in Fig. 2(b), the network given in Fig. 2(a) rapidly converges to a fixed point corresponding to the correct output for a given input, and the equilibrium state in each case may be taken as a code word C_0 or C_1 . The code words are shown overlayed on the network graphs in Fig. 2(c). The code words are of maximum Hamming distance from each other, meaning that C_0 differs from C_1 by every possible bit. This is desirable, as it means the two states are maximally distinguishable in the presence of noise. The evolved organism also stabilizes the code word: the state returns to the code word one or two time steps after perturbation by noise.

What is remarkable is that error correcting codes seem to be a typical result of this experiment. The lower half of Fig. 2 shows results from 150 independent evolutionary experiments. These populations were all completely randomly




FIG. 3. The effect of error correction on evolvability. (a) The lethality of different kinds of mutations for error correcting and non-error correcting AND organisms. A mutation is “lethal” if an organism that computes AND perfectly in the absence of noise can no longer do so after the mutation. Mutations to the truth tables consist of all possible bit flips applied to each node’s truth tables. Mutations to the connectivity constitute adding or deleting a single connection between nodes at random. The lethal mutation fraction is the fraction of all of the possible mutations that lead to death for a given organism. This is plotted as a distribution over all organisms for each error type. (b) Comparing the probability of surviving several sequential truth table mutations for error correcting and non-error correcting AND organisms. Error correcting organisms can survive three sequential mutations with the same probability that a non-error correcting organism can survive one. (c) Comparing the median logical error probability to the fraction of lethal truth table mutations for the populations of AND organisms from Fig. 2. We see that the lethal mutation fraction decreases with logical error, a direct demonstration of the principle of error correction enhanced evolvability.

initialized, and the results were not pruned: every population that was initialized has been included in the statistics. The distribution of population fitnesses over time is summarized in Fig. 2(c). As can be seen in Fig. 2(d), over 80% of the organisms compute the primitive with $p_{\log} < p_{\text{phys}}$, with the median organism suppressing error by nearly an order of magnitude. Figure 2(e) shows that organisms accomplish this using error correcting codes. Most organisms develop code words of nearly maximum Hamming distance. Figure 2(f) shows that the organism suppresses logical error superlinearly, a hallmark of error correction.

Achieving $p_{\log} < p_{\text{phys}}$ is an unambiguous symptom of fault tolerance: organisms that compute in a fragile way where a single error can lead to a logical fault will generally have $p_{\log} \gg p_{\text{phys}}$. Fragile networks experience such significant logical errors because they have many nodes and are exposed to noise for many iterations of the dynamics; see SM, Sec. 3 [45] for further discussion.

It should be emphasized here that no prior information about error correcting codes was supplied to the evolutionary procedure in the form of extra loss terms or otherwise: the mathematical concept of error correction emerges here exclusively from the primitive concepts of life and death in a noisy environment. This separates this result from previous studies where tailored optimization procedures were used explicitly to find good error correcting codes [58–62].

Another pattern present in the results presented in Fig. 2 is that the edges of the evolved Boolean networks are dominantly of moderate influence. The influence of an edge is the probability (taken over the entire truth table of the head node) that the state of the head node flips if the tail node were to flip. Edges of moderate influence are expected for a fault-tolerant solution, as heavy influence edges more readily propagate errors.

Interestingly, many Boolean network models of biological systems exhibit dynamics dominated by the effects of canalizing functions [63], which commonly involve edges with an influence close to 1. The presence of canalizing functions is often associated with stability, in the sense that the network will *eventually* return to its original attractor after a small perturbation. This notion of error resilience would not be sufficient for an organism to be considered fit in our experiments, which required errors to be corrected rapidly. We hypothesize that this difference in recovery timescales is the main reason our evolutionary process preferred distributed computation over canalizing functions.

The overwhelming typicality of error correcting codes in these results is surprising. Complete convergence of so many randomly initialized populations is not expected for adaptive walks on rugged loss landscapes [36]. We would expect many populations to get trapped in local minima and not find good codes. The implication of all of this is that in the presence of noise evolutionary processes somehow seek error correction. Note that the search of configuration space performed by evolution was nowhere near exhaustive. From Fig. 2(c), we can see that convergence was achieved in $\sim 10^6$ generations, with $\sim 10^2$ variations being explored each generation. The adaptive procedure explored a minuscule fraction of configuration space, $10^8/10^{40} = 10^{-32}$. This is far beyond finding a needle in a haystack ($\sim 10^{-10}$) and is closer to looking in the same haystack for a particular carbon atom.

One explanation is that error correcting codes not only protect an organism from noise that strikes randomly during the dynamics but also from *systematic* manipulation of the genome. This is supported by data presented in Fig. 3. For non-error correcting organisms, most genome mutations are lethal, such that an organism that perfectly performs a particular computation in the absence of noise before mutation no




FIG. 4. Emergent structure in noisy adapting dynamic systems. (a) Organisms that have adapted to solve the parallel and task with and without noise, all other experimental parameters held equal. The structures of the noisy organisms reflect the structure of the computational task, while the noiseless organisms have no obvious structure. This implies that noise alone can lead to modular dynamic systems. Experiments were also completed using the shared (b) and sequential (c) AND tasks. The computational structure of these tasks is also reflected beautifully in the evolved organisms. (d) An example of the code words developed for the shared AND task. The organism possesses two distinct memory modules of equal size, each of which implements an independent maximum Hamming distance code. The coloring of the nodes (white or black) indicates the steady-state value of the node (1 or 0) after the computation has taken place. (e) The code words associated with the shared AND task, which has three possible outputs. The code words are all of maximum Hamming distance from one another. (f) The structure of a network that computes a six-bit function before and after adaptation. We can see that the adaptive process imprints the computational structure on the organism, and also leads to the pruning of many unnecessary nodes.

longer does so after mutation. However, as shown in Fig. 3(a), the same is not true for error correcting organisms, for which only a small percentage of mutations are lethal.

The implication of this is error correction converts many lethal mutations into *neutral* mutations, such that organisms possessing error correcting codes experience more neutral mutations than those without. This is widely recognized as a sign of an evolvable property [64,65]. It follows that error correcting organisms can accumulate longer strings of mutations before suffering a lethal mutation, as shown in Fig. 3(b).

As illustrated in Fig. 3(c), the lethal mutation fraction decreases with logical error probability. Therefore, having some error correction makes it easier to search for more error correction, which will lead to error correction being a typical result in noisy evolutionary processes. This is the essence of the principle of error correction enhanced evolvability. This principle explains the prevalence of error correction in our experimental results, and possibly the natural world. Interestingly, noise is required to bootstrap this process, as there is no evolutionary pressure to develop an initial, small error correcting code without it. Noise enhances evolutionary processes.

Mutations are unit distance movements in configuration space. Random sequential mutations trace out random walks in configuration space, and a lethal mutation occurs when this walk encounters a configuration with a fitness that is too low to be survivable. Therefore, these results may be equivalently understood as error correction locally smoothing over valleys in the fitness landscape being searched by evolution. Figure 3(a) shows the smoothing of the region that can be

reached by a single mutation and therefore may be interpreted as error correction reducing the magnitude of the “gradient” of the fitness with respect to the configuration (although this is not a true gradient, since the configuration space is discrete). Figure 3(b) shows the same information over long distances. This concept of error correction smoothing the local fitness landscape and allowing error corrected organisms to search further for improvements is illustrated by the cartoon in Fig. 1.

IV. EVOLVING MODULARITY

Figure 4 shows that our results extend to composite computational tasks. The most striking feature of these organisms is that when evolved in the presence of noise, the structure of the computation is strongly imprinted on the structure of the organism. This is shown clearly in Fig. 4(a). Here, a parallel task is used, in which two noninteracting and computations must be completed. The task can be completed by two independent modules, each containing a pair of inputs and their corresponding outputs. When evolved in a noisy environment, this is exactly the structure present in the organisms. Independent modules saturate Newman’s modularity metric [66]: this solution could not be any more modular. When evolved in the absence of noise (with all other experimental parameters held constant) the structure vanishes, and the solutions resemble random graphs. The story is the same for tasks with more complicated modular structures, as shown in Figs. 4(b) and 4(c). It seems that noise, and noise alone,

can be enough to drive the evolution of structural modularity in dynamic systems.

Although the idea of modularity emerging naturally to increase the robustness of systems to noise is old [67], experimental demonstrations of this principle are rare. Understanding the emergence of modularity in cognition is a recent topic of interest. Most methods that have been developed involve imposing a special cost or environmental constraint that encourages modularity, such as rapidly varying a learning objective in a modular fashion [56,68] and associating a cost with adding connections to a network [69]. Results on noise alone leading to modularity are sparse, and limited to small neural networks [70]. Therefore, our results on modularity taken alone are significant, as they demonstrate a long-believed concept within the physically relevant formalism of dynamic systems.

It should be noted that like the primitive organisms, the organisms shown in Fig. 4 also possess error correcting codes. As shown in Fig. 4(d), the structurally distinct modules seen in the shared and organism also make up functionally distinct memory modules. Each memory module implements an independent maximum Hamming distance code for each output bit, akin to the code shown in Fig. 1. This modular error correction reflects the way the brain seems to implement independent error correcting codes in different regions [18,19,23]. The way the modules are constructed is also interesting, as from a fitness standpoint each output bit was made equally important and the same number of nodes are dedicated to each module. This kind of code is formally efficient in the sense that the number of bits it can encode scales linearly with the number of physical bits (it is asymptotically a constant rate code). The sequential and organism develops a different kind of code, as shown in Fig. 4(e). In this case, memory modules would not be an efficient solution, as there are only three possible outputs for this task: 00, 01, and 11. Instead, this organism develops a Hamming-like [71] code that maximizes the distance between each of the three distinct code words.

To explore the emergence of modularity at larger scales while maintaining computational tractability, instead of starting searches from random states one can seek to improve the fault tolerance of a network that already performs a particular computation in the absence of noise. In particular, we sample random networks that compute a function on a certain number of bits in the absence of noise, and then evolve them in a noisy environment and observe the results. An example of an organism that implements a six-bit function is shown in Fig. 4(f). We can see that substantial structure is evolved. The final network is modular, with nodes 1 and 3 being distinctly separated from the other nodes. Additionally, the evolutionary process removes many of the nodes present in the original network, suggesting they were useless to the computation and only increased noise cross section.

V. DISCUSSION

It can be hypothesized that the general principle behind noise-induced modularity is that noise punishes network size, both in terms of the number of connections and the number of

nodes. Noise punishes connectivity because errors can propagate further and faster in more densely connected networks. For example, the separate modules evolved to solve the noisy parallel and task prevent errors that strike one module from ever propagating to the other. In general, connectivity in the organisms evolved to solve noisy tasks reflects the computationally required connectivity to move information between inputs and outputs. The shared and organisms generally separate into two modules, with connections present between the two only to communicate the state of node 1. Similar features are observed in the noisy sequential and organisms. Noise punishes network size because larger networks have more components subject to error.

The experimental results presented in this paper demonstrate that error correction and modularity go hand in hand, and may be expected to emerge as a result of noisy evolutionary processes. Explicitly, we show that error correction is a typical result of evolving to perform a primitive computation in the presence of noise and that this principle scales to small composite computations, where solutions that are both modular and error correcting emerge.

We reason that this occurs because of error correction enhanced evolvability: systems possessing error correcting codes are easier to improve.

Note that it has previously been hypothesized that robustness in general may be evolvable [72]. Here, we provide specific experimental evidence that backs this claim, and find a particular operational mechanism: the gradual accumulation of an error correcting code and the associated fault-tolerance machinery.

With our principle in hand, we believe that these results should extend to scales more comparable in size to biological information processing systems. Such regimes are out of reach of our current techniques as evolving orders of magnitude larger organisms is presently computationally intractable. Future work may explore alternative modeling mechanisms, including direct physical implementations of noisy dynamics with electrical, mechanical, or biological mechanisms, or with neural network learning as a proxy for evolutionary change.

ACKNOWLEDGMENTS

This work was supported in part by the Institute for Artificial Intelligence and Fundamental Interactions (IAIFI) through NSF Grant No. PHY-2019786, and by NTT Research. T.M. would like to thank Cian Schmitt-Ulms, Olive Garst, and Mikail Khona for their useful comments that helped clarify the manuscript. T.M. would also like to thank the administrators of the subMIT cluster, on which many of the experiments presented in this work were completed.

T.M. designed and ran the experiments, and analyzed the results. T.M., I.R.F., and I.L.C. wrote and edited the manuscript. T.M. wrote the Supplemental Material.

DATA AVAILABILITY

The data that support the findings of this article are openly available [44,55].

- [1] A. Einstein, On the motion of small particles suspended in a stationary liquid, *Ann. Phys.* **322**, 549 (1905).
- [2] J. von Neumann, The general and logical theory of automata, in *Cerebral Mechanisms in Behaviour*, edited by L. A. Jeffress (Wiley, Hoboken, NJ, 1951).
- [3] J. V. Neumann and A. W. Burks, *Theory of Self-Reproducing Automata* (University of Illinois Press, Champaign, ILU, 1966).
- [4] J. von Neumann, Probabilistic logics and the synthesis of reliable organisms from unreliable components, in *Automata Studies (AM-34), Vol. 34*, edited by C. E. Shannon and J. McCarthy (Princeton University Press, Princeton, NJ, 1956), pp. 43–98.
- [5] W. C. Carter, Fault-tolerant computing: An introduction and a viewpoint, *IEEE Trans. Comput.* **C-22**, 225 (1973).
- [6] A. Zlokap, A. K. Tan, J. M. Martyn, I. R. Fiete, M. Tegmark, and I. L. Chuang, Fault-tolerant neural networks from biological error correction codes, *Phys. Rev. E* **110**, 054303 (2024).
- [7] P. W. Shor, Fault-tolerant quantum computation, in *Proceedings of the 37th Conference on Foundations of Computer Science* (IEEE, Piscataway, NJ, 1996), p. 56.
- [8] G. Q. Ai and S. Feynman, Suppressing quantum errors by scaling a surface code logical qubit, *Nature (London)* **614**, 676 (2023).
- [9] C. K. Andersen, A. Remm, S. Lazar, S. Krinner, N. Lacroix, G. J. Norris, M. Gabureac, C. Eichler, and A. Wallraff, Repeated quantum error detection in a surface code, *Nat. Phys.* **16**, 875 (2020).
- [10] S. Krinner, N. Lacroix, A. Remm, A. D. Paolo, E. Genois, C. Leroux, C. Hellings, S. Lazar, F. Swiadek, J. Herrmann, G. J. Norris, C. K. Andersen, M. Müller, A. Blais, C. Eichler, and A. Wallraff, Realizing repeated quantum error correction in a distance-three surface code, *Nature (London)* **605**, 669 (2022).
- [11] C. E. Shannon, A mathematical theory of communication, *Bell Syst. Tech. J.* **27**, 379 (1948).
- [12] S. Dow, A. Lucan-danila, K. Anderson, A. P. Arkin, A. Astromoff, M. E. Bakkoury, R. Bangham, R. Benito, S. Brachat, B. Andre, D. F. Jaramillo, D. E. Kelly, S. L. Kelly, and P. Ko, Functional profiling of the *Saccharomyces cerevisiae* genome, *Nature (London)* **418**, 387 (2002).
- [13] D. Deutscher, I. Meilijson, M. Kupiec, and E. Rupp, Multiple knockout analysis of genetic robustness in the yeast metabolic network, *Nat. Genet.* **38**, 993 (2006).
- [14] D. A. Drachman, Do we have brain to spare? *Neurology* **64**, 2004 (2005).
- [15] T. H. Mosley, D. S. Knopman, D. J. Catellier, N. Bryan, R. G. Hutchinson, C. A. Grothues, A. R. Folsom, L. S. Cooper, G. L. Burke, D. Liao, and M. Szklo, Cerebral MRI findings and cognitive functioning: The Atherosclerosis Risk in Communities Study, *Neurology* **64**, 2056 (2005).
- [16] L. C. B. Faria, A. S. L. Rocha, J. H. Kleinschmidt, M. C. Silva-Filho, E. Bim, R. H. Herai, M. E. B. Yamagishi, and R. Palazzo, Is a genome a codeword of an error-correcting code? *PLoS One* **7**, e36644 (2012).
- [17] K. Yoon, M. A. Buice, C. Barry, R. Hayman, N. Burgess, and I. R. Fiete, Specific evidence of low-dimensional continuous attractor dynamics in grid cells, *Nat. Neurosci.* **16**, 1077 (2013).
- [18] R. Chaudhuri, B. Gerçek, B. Pandey, A. Peyrache, and I. Fiete, The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep, *Nat. Neurosci.* **22**, 1512 (2019).
- [19] R. J. Gardner, E. Hermansen, M. Pachitariu, Y. Burak, N. A. Baas, B. A. Dunn, M. B. Moser, and E. I. Moser, Toroidal topology of population activity in grid cells, *Nature (London)* **602**, 123 (2022).
- [20] M. Khona and I. R. Fiete, Attractor and integrator networks in the brain, *Nat. Rev. Neurosci.* **23**, 744 (2022).
- [21] J. D. Franson, Quantum computing: A topological route to error correction, *Nature* **482**, 479 (2012).
- [22] *Quantum Error Correction*, edited by D. A. Lidar and T. A. Brun (Cambridge University Press, Cambridge, 2013).
- [23] S. Sreenivasan and I. Fiete, Grid cells generate an analog error-correcting code for singularly precise neural computation, *Nat. Neurosci.* **14**, 1330 (2011).
- [24] L. Zheng and W. Wang, Regulation associated modules reflect 3D genome modularity associated with chromatin activity, *Nat. Commun.* **13**, 5281 (2022).
- [25] D. Meunier, R. Lambiotte, and E. T. Bullmore, Modular and hierarchically modular organization of brain networks, *Front. Neurosci.* **4** (2010).
- [26] M. Landau, Redundancy, rationality, and the problem of duplication and overlap, *Pub. Adm. Rev.* **29**, 346 (1969).
- [27] A. W. Lerner, There is more than one way to be redundant: A comparison of alternatives for the design and use of redundancy in organizations, *Adm. Soc.* **18**, 334 (1986),.
- [28] I. Nonaka, Redundant, overlapping organization: A Japanese approach to managing the innovation process, *Calif. Manage. Rev.* **32**, 27 (1990),.
- [29] M. Aldana, E. Balleza, S. Kauffman, and O. Resendiz, Robustness and evolvability in genetic regulatory networks, *J. Theor. Biol.* **245**, 433 (2007).
- [30] A. Wagner, Robustness and evolvability: A paradox resolved, *Proc. R. Soc. B* **275**, 91 (2008).
- [31] A. Kucharavy, B. Rubinstein, J. Zhu, and R. Li, Robustness and evolvability of heterogeneous cell populations, *Mol. Biol. Cell* **29**, 1400 (2018).
- [32] N. Arndt, Hot heads and cold tails, *Nature (London)* **407**, 459 (2000).
- [33] M. Pigliucci, Is evolvability evolvable? *Nat. Rev. Genet.* **9**, 75 (2008).
- [34] J. L. Payne and A. Wagner, The causes of evolvability and their evolution, *Nat. Rev. Genet.* **20**, 24 (2019).
- [35] E. D. Vaishnav, C. G. de Boer, J. Molinet, M. Yassour, L. Fan, X. Adiconis, D. A. Thompson, J. Z. Levin, F. A. Cubillos, and A. Regev, The evolution, evolvability and engineering of gene regulatory DNA, *Nature (London)* **603**, 455 (2022).
- [36] S. A. Kauffman, *The Origins of Order: Self-Organization and Selection in Evolution* (Oxford University Press, New York, 1993).
- [37] C. Torres-Sosa, S. Huang, and M. Aldana, Criticality is an emergent property of genetic networks that exhibit evolvability, *PLoS Comput. Biol.* **8**, e1002669 (2012).
- [38] S. Kauffman, Metabolic stability and epigenesis in randomly constructed genetic nets, *J. Theor. Biol.* **22**, 437 (1969).
- [39] S. A. Kauffman, Homeostasis and differentiation in random genetic control networks, *Nature (London)* **224**, 177 (1969).
- [40] R. Thomas, Boolean formalization of genetic control circuits, *J. Theor. Biol.* **42**, 563 (1973).
- [41] T. Parr, G. Pezzulo, and K. J. Friston, *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior* (The MIT Press, Cambridge, MA, 2022).

- [42] S. Kauffman, Antichaos and adaptation, *Sci. Am.* **265**, 78 (1991).
- [43] F. Bertacchini, C. Scuro, P. Pantano, and E. Bilotta, Modelling brain dynamics by Boolean networks, *Sci. Rep.* **12**, 16543 (2022).
- [44] T. McCourt, GPU accelerated Kauffman networks (2023), https://github.com/trevormccrt/gpuaffman_networks.
- [45] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/7kcl-jshm> for additional derivations and simulations that provide extra support for our results, which also includes Refs. [73–76].
- [46] C. Gershenson, Introduction to random Boolean networks, [arXiv:nlin/0408006](https://arxiv.org/abs/nlin/0408006).
- [47] I. Harvey and T. Bossomaier, Time out of joint: Attractors in asynchronous random Boolean networks, in *Proceedings of the Fourth European Conference on Artificial Life*, edited by P. Husbands and I. Harvey (The MIT Press, Cambridge, MA, 1997), p. 67.
- [48] M. M. Vopson, Estimation of the information contained in the visible matter of the universe, *AIP Adv.* **11**, 105317 (2021).
- [49] I. Shmulevich, E. R. Dougherty, S. Kim, and W. Zhang, Probabilistic Boolean networks: A rule-based uncertainty model for gene regulatory networks, *Bioinformatics* **18**, 261 (2002).
- [50] C. Ofria and C. O. Wilke, Avida: A software platform for research in computational evolutionary biology, *Artif. Life* **10**, 191 (2004).
- [51] S. S. Chow, C. O. Wilke, C. Ofria, R. E. Lenski, and C. Adami, Adaptive radiation from resource competition in digital organisms, *Science* **305**, 84 (2004).
- [52] M. A. Fortuna, L. Zaman, A. P. Wagner, and C. Ofria, Evolving digital ecological networks, *PLoS Comput. Biol.* **9**, e1002928 (2013).
- [53] J. H. Holland, *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence* (The MIT press, Cambridge, MA, 1992).
- [54] C. Adami, C. Ofria, and T. C. Collier, Evolution of biological complexity, *Proc. Natl. Acad. Sci. USA* **97**, 4463 (2000).
- [55] T. McCourt, Code to accompany the paper “emergent efficient error correction and modularity in noisy dynamic systems” (2023), https://github.com/trevormccrt/noise_emergence.
- [56] N. Kashtan and U. Alon, Spontaneous evolution of modularity and network motifs, *Proc. Natl. Acad. Sci. USA* **102**, 13773 (2005).
- [57] J. P. Crutchfield and M. Mitchell, The evolution of emergent computation. *Proc. Natl. Acad. Sci. USA* **92**, 10742 (1995).
- [58] A. Elkelesh, M. Ebada, S. Cammerer, and S. t. Brink, Decoder-tailored polar code design using the genetic algorithm, *IEEE Trans. Commun.* **67**, 4521 (2019).
- [59] L. Huang, H. Zhang, R. Li, Y. Ge, and J. Wang, AI coding: Learning to construct error correction codes, *IEEE Trans. Commun.* **68**, 26 (2020).
- [60] M. D. Jaraiz Simon, J. A. Gomez Pulido, M. A. Vega Rodriguez, J. M. Sanchez Perez, and J. M. Granado Criado, A genetic algorithm to design error correcting codes, in *MELECON 2006–2006 IEEE Mediterranean Electrotechnical Conference* (IEEE, Piscataway, NJ, 2006), pp. 807–810.
- [61] H. Kim, Y. Jiang, R. B. Rana, S. Kannan, S. Oh, and P. Viswanath, Communication algorithms via deep learning, in *International Conference on Learning Representations* (ICLR, 2018).
- [62] H. Kim, Y. Jiang, S. Kannan, S. Oh, and P. Viswanath, Deep-Code: Feedback codes via deep learning, in *Advances in Neural Information Processing Systems 31 (NeurIPS 2018)*, edited by S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, R. Garnett (Curran Associates, Red Hook, NY, 2018).
- [63] S. Kauffman, C. Peterson, B. Samuelsson, and C. Troein, Genetic networks with canalizing Boolean rules are always stable, *Proc. Natl. Acad. Sci. USA* **101**, 17102 (2004).
- [64] O. Tenaillon and I. Matic, The impact of neutral mutations on genome evolvability, *Curr. Biol.* **30**, R527 (2020).
- [65] R. E. Lenski, J. E. Barrick, and C. Ofria, Balancing robustness and evolvability, *PLoS Biol.* **4**, e428 (2006).
- [66] M. E. J. Newman, Modularity and community structure in networks, *Proc. Natl. Acad. Sci. USA* **103**, 8577 (2006).
- [67] G. Wagner, M. Pavlicev, and J. Cheverud, The road to modularity, *Nat. Rev. Genet.* **8**, 921 (2007).
- [68] N. Kashtan, E. Noor, and U. Alon, Varying environments can speed up evolution, *Proc. Natl. Acad. Sci. USA* **104**, 13711 (2007).
- [69] J. Clune, J. B. Mouret, and H. Lipson, The evolutionary origins of modularity, *Proc. R. Soc. B* **280**, 20122863 (2013).
- [70] B. A. Høverstad, Noise and the evolution of neural network modularity, *Artif. Life* **17**, 33 (2011).
- [71] R. W. Hamming, Error detecting and error correcting codes, *Bell Syst. Tech. J.* **29**, 147 (1950).
- [72] H. Kitano, Biological robustness, *Nat. Rev. Genet.* **5**, 826 (2004).
- [73] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, *et al.*, Array programming with NumPy, *Nature (London)* **585**, 357 (2020).
- [74] R. Okuta, Y. Unno, D. Nishino, S. Hido, and C. Loomis, CuPy: A NumPy-compatible library for NVIDIA GPU calculations, in *Proceedings of Workshop on Machine Learning Systems (LearningSys) in The Thirty-first Annual Conference on Neural Information Processing Systems (NIPS)* (Curran Associates, Red Hook, NY, 2017).
- [75] M. Minsky and S. A. Papert, *Perceptrons: An Introduction to Computational Geometry* (The MIT Press, Cambridge, MA, 2017).
- [76] A. Globus, E. Langhirt, M. Livny, R. Ramamurthy, M. Solomon, and S. Traugott, JavaGenes and Condor: Cycle-scavenging genetic algorithms, in *Proceedings of the ACM Conference on Java Grande* (ACM, New York, NY, 2002), pp. 134–139.