Noisy dynamical systems evolve error correcting codes and modularity *

Trevor McCourt[1], Ila R. Fiete[1] and Isaac L. Chuang[1]

[1]Massachusetts Institute of Technology, Cambridge, MA 02139, USA

Even under the influence of entirely deterministic natural laws, an agent with only finite information about its surroundings will experience some degree of randomness. This is exemplified in physics by observations of Brownian motion Einstein (1905), 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 that emerges, it is doing so despite a stochastic environment, as it has on Earth.

At the same time, living may be phrased as computation. Receiving a stimulus from the environment and responding in the "correct" way (that maximizes one's ability to continue to stay alive) is equivalent to computing the output of a function given an input Parr et al. (2022). Given this and the ubiquity of noise, it is clear that there is an intimate connection between life and fault-tolerant computing, as first recognized by von Neumann. Inspired by the incredible degree of robust computing he observed in the natural world, von Neumann developed schemes for robustly performing computations using only noisy components von Neumann (1956).

Since the time of von Neumann, we have made substantial experimental progress toward understanding the mechanisms biological systems use to compute fault tolerantly. Namely, we have discovered many examples of formal error-correcting codes Shannon (1948) in biological systems. Crick famously discovered that a simple form of a triple redundancy code is used to reliably synthesize proteins from mRNA Crick et al. (1957). 21^{st} century geneticists have since discovered more explicit examples of error correction: it has been found that DNA sequences, and even entire genomes, can be identified as codewords of the famous Hamming code Hamming (1950); Faria et al. (2012). Error-correcting codes are also a hot topic in neuroscience. In particular, several examples of low dimensional manifolds have been discovered in the activity of particular neural circuits Chaudhuri et al. (2019); Gardner et al. (2022), which can be identified as topological codes. In particular, we have discovered that grid cells, which are part of the circuitry that help us reason spatially, implement an intricate error-correcting code called the grid code Sreenivasan and Fiete (2011).

Biological systems also often feature a high degree of modularity, which is the physical division of a large system into functionally distinct parts. This is thought to contribute to their robustness, as the failure of one subsystem will not necessarily cascade and cause the failures of others. The genome is believed to have spatial and functional modularity Zheng and Wang (2022), 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 Meunier et al. (2010).

On the other hand, all biological systems have emerged through the brutally simple process of evolution: improving fitness via essentially random genetic mutations. It is incredible that evolution was able to produce modular, error-correcting structures through this process, given that they represent a vanishing fraction of the configuration space of possible solutions to a given problem. A tantalizing question is therefore what conditions and principles drive evolution toward such structures? The answer to this question is significant, as it gets to the root of forces that drive the crystallization of life.

In this work, via experiments in Boolean networks, we show that both error-correcting codes and modularity are typical co-emergent 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. Noise bootstraps this phenomenon, suggesting that noise plays an important role in evolving complex structures.

Our method is to explicitly exploit the connection between fault-tolerant computing and life to understand how error correction and modularity may emerge via evolution.

*This is an extended abstract prepared for an ALIFE 2023 workshop. The full manuscript can be found at arXiv:2303.14448

We use Boolean networks as examples of primitive artificial lifeforms and evaluate them based on their ability to perform a computation in the presence of noise. Explicitly, as shown in Fig. 1 a this is done by initializing a network with the inputs to a computational task via assigned input nodes (orange) and probing it for the output by reading from a set of output nodes (green) after many dynamical timesteps. Noise is injected during the dynamics by randomly flipping bits of the network state with probability p_{phys}. The logical error probability p_{log} is the probability that the output node does not contain the correct value at the end of the computation. For a given p_{phys}, networks with lower p_{log} are more effectively performing the computation in the presence of noise. Starting from completely random initializations, we use p_{log} as a measure of fitness in an evolutionary procedure to generate improved networks. Fig. 1 b shows the dynamical trajectory of the network shown in Fig. 1 a, which was evolved to compute AND. The network has evolved substantial fault tolerance: it robustly performs the AND computation, encodes the outputs in codewords of maximum Hamming distance (the codewords differ by every bit), and stabilizes them against noise. Evolution has discovered a strong error correction mechanism! Remarkably, contrary to intuition on optimization on massive, rugged fitness landscapes, evolution seems to routinely find these codes. Fig. 1 c shows the evolutionary trajectory of several hundred randomly initialized populations, almost all of which achieve $p_{\text{phys}} < p_{\text{log}}$, as shown in Fig. 1 d. Fig. 1 e shows that networks achieve this by implementing strong error-correcting codes like the one in Fig. 1 b, which allows p_{log} to scale less than linearly with p_{phys}, as shown in Fig. 1 f.

One explanation for this overwhelming typicality of error correction is that error-correcting codes not only protect an organism from the noise that strikes randomly during the dynamics but also from *systematic* manipulation of the genome. This is supported by data presented in Fig. 1 g. We can see that as error correction improves (and p_{log} decreases), there is a corresponding decrease in the number of mutations that are lethal to an organism (where a lethal mutation is one that prevents an organism from completing the desired computation perfectly in the absence of noise). Therefore, acquiring some error correction makes it easier to acquire more error correction, leading to an "explosion of error correction" in our experimental results, and possibly the natural world. We dub this principle error correction enhanced evolvability.

Additionally, Figs. 1 h and i show that when scaled to composite computational tasks (that require more than one primitive), the structure of the computation is firmly imprinted on the structure of the organism. Noise promotes modularity in dynamic systems. These tasks also have output spaces larger than 1 bit, and demand more advanced error-correcting codes than what is shown in Fig. 1 b. The shared AND task, as shown in Fig. 1 h, has a 2-bit output space. Interestingly, independent memory modules are developed that correspond to the physical modules, such that one bit of the output is coded for in 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). This modular assembly of codes also reflects beautifully what we see in the brain. The sequential AND task, as shown in Fig. 1 i has 3 possible outputs (less than two bits), so it develops a distributed code where each codeword is of maximum hamming distance.

References

Chaudhuri, R., Gerçek, B., Pandey, B., Peyrache, A., and Fiete, I. (2019). The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nature Neuroscience*, 22(9):1512–1520.

Crick, F. H. C., Griffith, J. S., and Orgel, L. E. (1957). Codes without commas. *Proceedings of the National Academy of Sciences*, 43(5):416–421.

Einstein, A. (1905). On the motion of small particles suspended in a stationary liquid. *Ann. Phys*, 322(8):549–560.

Faria, L. C., Rocha, A. S., Kleinschmidt, J. H., Silva-Filho, M. C., Bim, E., Herai, R. H., Yamagishi, M. E., and Palazzo, R. (2012). Is a genome a codeword of an error-correcting code? *PLoS ONE*, 7(5).

Gardner, R. J., Hermansen, E., Pachitariu, M., Burak, Y., Baas, N. A., Dunn, B. A., Moser, M. B., and Moser, E. I. (2022). Toroidal topology of population activity in grid cells. *Nature*, 602(7895):123–128.

Hamming, R. W. (1950). Error detecting and error correcting codes. *The Bell System Technical Journal*, 29(2):147–160.

Meunier, D., Lambiotte, R., and Bullmore, E. (2010). Modular and hierarchically modular organization of brain networks. *Frontiers in Neuroscience*, 4.

Parr, T., Pezzulo, G., and Friston, K. J. (2022). *Active Inference: The Free Energy Principle in Mind, Brain, and Behavior*. The MIT Press.

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423.

Sreenivasan, S. and Fiete, I. (2011). Grid cells generate an analog error-correcting code for singularly precise neural computation. *Nature Neuroscience*, 14(10):1330–1337.

von Neumann, J. (1956). *Probabilistic Logics and the Synthesis of Reliable Organisms From Unreliable Components*, pages 43–98. Princeton University Press, Princeton.

Zheng, L. and Wang, W. (2022). Regulation associated modules reflect 3d genome modularity associated with chromatin activity. *Nature Communications*, 13.

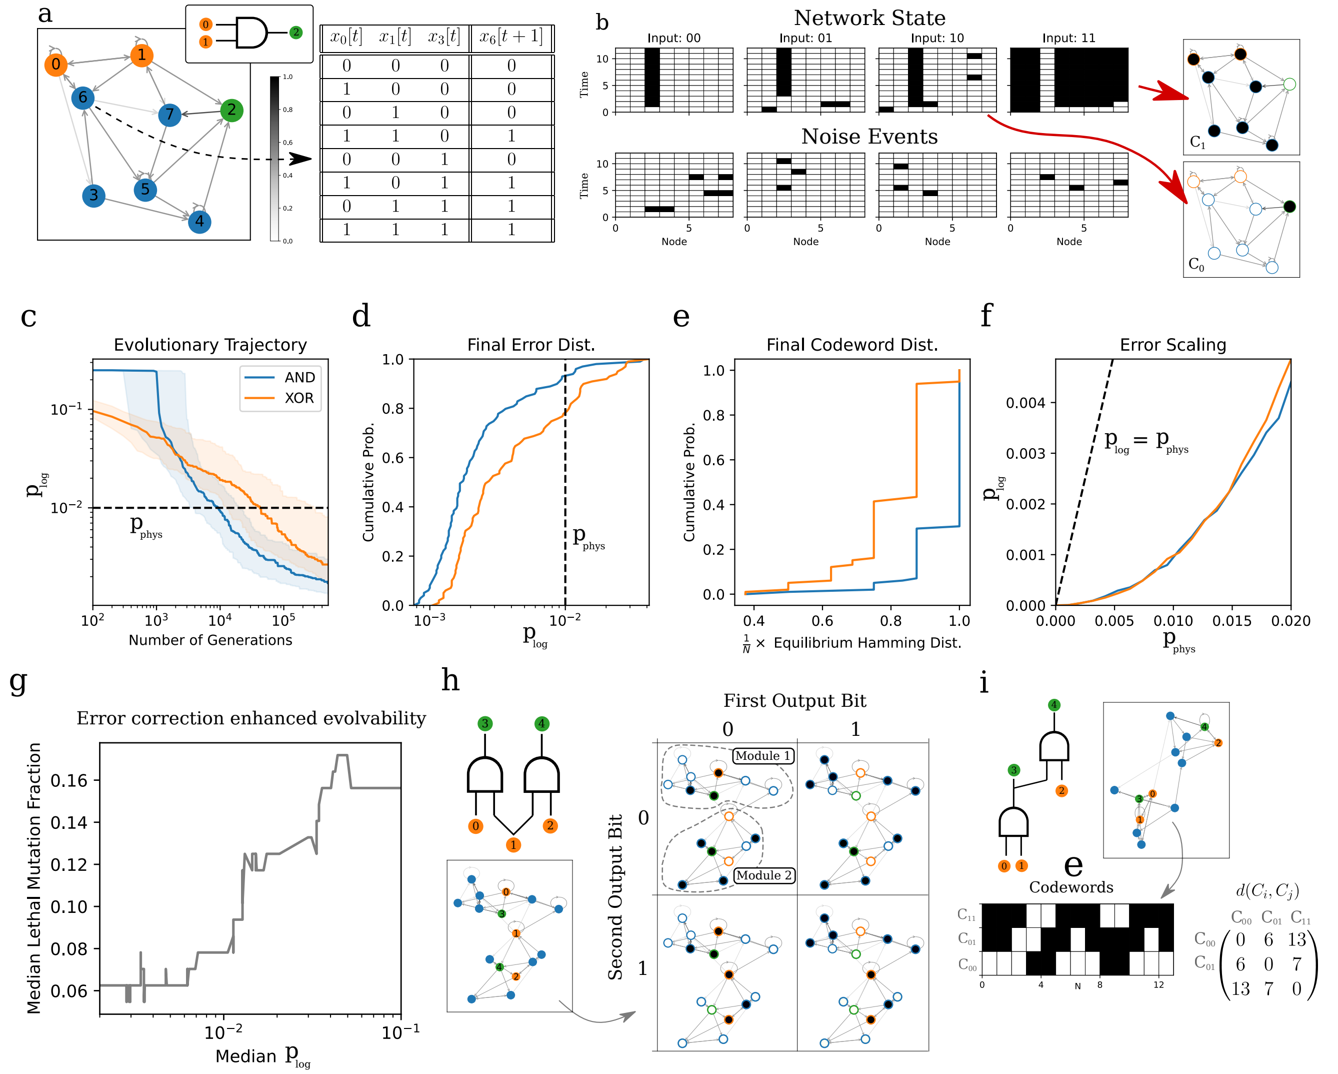

Figure 1: **Noisy Boolean networks evolve strong error-correcting codes and modularity a** A Boolean network that has adapted to solve the AND task in the presence of noise. Inputs are provided to the orange nodes 0 and 1 at time t=0 and, the answer is expected at the green node 2 at t=T. **b** Examples of noisy dynamical trajectories of the network in (a) for each of the 2^2 possible input states. The network encodes the answer in codewords of maximal Hamming distance and stabilizes the codeword against noise events. Codewords may be taken as the final state of the network $x[T]$ for the different output values. The two codewords are visualized on the network graph. The boundary of the circles indicates the node function, and the fill indicates the codeword. **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 codewords 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. **g** Comparing the median logical error probability to the fraction of lethal truth table mutations for the populations of AND organisms from (c). We see that the lethal mutation fraction decreases with logical error, a direct demonstration of the principle of error correction enhanced evolvability. **h** Organisms evolved to solve the shared AND task develop a strikingly modular physical structure and also employ a modular error correcting code, where each output bit is independently encoded by the nodes in the corresponding module. **I** The sequential AND task also develops modular structure, but not modular error correction. It instead develops a distributed, maximum Hamming distance code.