Circuit design in biology and machine learning. I. Random networks and dimensional reduction

Steven A. Frank*

Department of Ecology and Evolutionary Biology, University of California, Irvine, CA, United States *Corresponding author: Department of Ecology and Evolutionary Biology, University of California, Irvine, CA 92697–2525, United States. Email: saf@stevefrank.org

Abstract

A biological circuit is a neural or biochemical cascade, taking inputs and producing outputs. How have biological circuits learned to solve environmental challenges over the history of life? The answer certainly follows Dobzhansky's famous quote that "nothing in biology makes sense except in the light of evolution." But that quote leaves out the mechanistic basis by which natural selection's trial-and-error learning happens, which is exactly what we have to understand. How does the learning process that designs biological circuits actually work? How much insight can we gain about the form and function of biological circuits by studying the processes that have made those circuits? Because life's circuits must often solve the same problems as those faced by machine learning designs computational circuits to solve problems. We can then ask: How much insight do those computational circuits provide about the design of biological circuits? How much does biology differ from computers in the particular circuit designs that it uses to solve problems? This article steps through two classic machine learning models to set the foundation for analyzing broad questions about the design of biological circuits. One insight is the surprising power of randomly connected networks. Another is the central role of internal models of the environment embedded within biological circuits, illustrated by a model of dimensional reduction and trend prediction. Overall, many challenges in biology have machine learning analogs, suggesting hypotheses about how biology's circuits are designed.

Keywords: biological design, evolution, artificial intelligence, reservoir networks, recurrent neural networks

Introduction

Biological circuits process information and make decisions. For example, a metabolic response circuit may encode the recent availability of various nutrients into a low-dimensional internal representation within the cell. The circuit then decodes the internal state into an appropriate metabolic response (Wu et al., 2022).

Similarly, machine learning circuits also process information and make decisions. Particular computational circuits reduce dimension, classify inputs, predict environments, and respond accordingly (Goodfellow et al., 2016).

Natural and artificial systems design their circuits by learning. Biology uses natural selection's trial-and-error learning. Machines use various learning algorithms. In each case, credit for performance feeds back to circuit components. Update rules reshape the circuits to improve future performance.

How similar are biology's and machine learning's circuits? Mostly, we do not know. I review analogies and highlight topics for study.

Importantly, what we see in machines suggests what to look for in biology. And what works in biology provides insight into machines. We gain a broader perspective on generalized notions of learning and the circuit architecture of learned designs (Aldarondo et al., 2024; Kanwisher et al., 2023). What about intelligence? Some people have argued that machine learning exploits correlations, whereas biological intelligence uses a causal model of the world (Mitchell, 2023). But what if machines can discover and use causal models? In any case, the comparison between machines and biology forces clearer thinking about how learning shapes response circuit design.

Induction by comparative example

The distinct materials of machines and biology inevitably cause differences in speed, efficiency, work, and cost. But, for relatively simple challenges, I do not know of limitations on information processing that require machines and organisms to solve problems with significantly different circuit designs.

If anything, biology's slow chemistry and evolutionary algorithms for learning seem more constrained than machine learning's fast electricity and diverse algorithms, whereas biology's intrinsic parallelism and energy efficiency pose some engineering challenges for machines. These material differences must cause low-level mechanistic differences in response circuits. But we do not know whether high-level structural and functional differences often arise by necessity (Baltussen et al., 2024).

Received August 19, 2024; revisions received November 13, 2024; accepted March 27, 2025 Handling Editor: Mark Kirkpatrick

[©] The Author(s) 2025. Published by Oxford University Press on behalf of The Society for the Study of Evolution (SSE).

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (https://creativecommons.org/licenses/by/4.0/),

which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

Put another way, we lack well-developed theory that predicts essential differences between the domains. We must proceed inductively by accumulating various comparisons between machine learning and biology. How do the distinct systems solve particular challenges? Architecturally, what sorts of similarities and differences do we see in the design of response circuits? Are there qualitative differences in information flow and decision making?

Ultimately, what sorts of predictions can we make about the design of biological response circuits? What sorts of inductive generalities arise about how learning shapes circuits and systems?

To start on these questions, I begin with some basic machine learning circuits. I discuss possible links between each computational example and biological circuits. I emphasize simple genetic and biochemical circuits in cells, for which there are relatively few prior studies of this sort. I occasionally mention neurobiological networks, for which there are many previous analyses that link computational and biological circuits (Hebb, 1949; Jiahui et al., 2023; Kanwisher et al., 2023; Miikkulainen, 2025; McCulloch & Pitts, 1943; Rumelhart & McClelland, 1986).

Overview

Each article in this series develops particular comparisons between machines and biology. This first article starts with randomly connected reservoir networks in machine learning. Such reservoir networks can be remarkably successful for certain types of challenges. For example, a randomly connected circuit can often predict future inputs given the recent sequence of input values.

Reservoir networks introduce several general principles of circuit design. The following section summarizes those principles. After that, I consider possible biological insights from the study of random networks in machine learning.

One insight concerns the evolutionary origin of specific circuit designs. Consider perception and response. Perception provides little value without a matching response. How do things get started? Machine learning's random reservoir networks suggest a path.

Random networks also illustrate how circuits gain from reducing the dimensionality of their inputs. The second part of this article continues analysis of dimensional reduction, a process that benefits circuit designs for many different kinds of challenges.

For example, recognizing an ear as a facial component in an image is the reduction of a high-dimension array of pixels to a lower-dimension model of a face. Similarly, a biological cell may reduce the abundances of various nutrients to a single dimensional measure of beneficial metabolic state. Machine learning models provide many insights into how simple circuit designs can reduce dimension to solve challenges.

The next article in this series discusses anomaly detection. In machine learning, a particular challenge might be detecting a rare fraudulent credit card transaction among a vast number of legitimate transactions. What pattern of inputs suggests that something is atypical? The problem is easy if certain features associate with danger. The problem is hard if anomalies often arise in different ways so that all one can do is identify anomalies by an atypical multivariate pattern.

A similar challenge arises in biology. Many particular signs of danger can be detected relatively easily. But how does an organism identify trouble by sensing a nonspecific deviation from the typical multivariate pattern of inputs? The methods and examples from machine learning provide general insights into the circuit designs that can solve this challenge.

For example, in machine learning, what characteristics and patterns of sensor arrays help in detecting anomalies? How does combining information from different circuits in an ensemble improve success? What is the role of sequential improvement in circuit performance, an analogy with the evolutionary dynamics of circuit improvement in biology? What hypotheses do we gain about how simple biological circuits may be designed?

Future articles in the series will explore such questions. Topics include how circuits manage uncertainty, how they learn to embed models of the world, and how machine learning techniques aid in designing simple biochemical circuits. Another topic examines how robustness influences circuit design and complexity.

Reservoir computing

This section illustrates how randomly connected networks store information about recent inputs. That reservoir of information about the past can be used to predict future environmental states and to respond accordingly. Random reservoir networks first arose in computational modeling (Cucchi et al., 2022; Gauthier et al., 2021; Jaeger, 2007; Maass et al., 2002). Several studies then considered how biology exploits similar reservoir networks (Czégel et al., 2021; Damicelli et al., 2022; Goudarzi et al., 2013; Loeffler et al., 2021; Seoane, 2019; Solé & Seoane, 2022; Yahiro et al., 2018). Box 1 and Fig. 1 provide background on reservoir computing.

Box 1: Reservoir computing

Reservoir computing (RC) provides a model for how complex computations can emerge from seemingly simple network structures, mirroring the distributed processing observed in biological circuits.

The core of RC is the reservoir, a network of interconnected nodes, much like a network of neurons or a group of interacting molecules in a biochemical network of reactions. These nodes are linked by connections that allow information to flow between them. The connectivity is **recurrent**, meaning that information can flow in feedback loops. For example, node A can influence node B, which can influence node C, which in turn can loop back and influence node A again.

This circular flow of information, a continuous feedback loop, allows the system to maintain a memory of past signals and process information over time. Think of it as a dynamic echo chamber where incoming signals reverberate and interact, creating a complex internal representation of the input history.

Importantly, these internal connections are mostly randomly formed and static. They do not change during computation, and they are not improved or optimized as the system learns or evolves. Instead, they simply provide a passive vessel that stores a trace of past inputs. The inputs can be external environmental signals or internal signals arising from other components of the system.

Neuronal and gene regulatory networks also have recurrent connectivity, which plays an important role in processing temporal information. The random connectivity of reservoirs and the lack of change in connectivity over time suggest that precise, predetermined wiring is not necessary for memory or complex computation.

RC has several properties that may be important in biology. **Temporal memory** arises through recurrent connections, enabling the reservoir to maintain a fading memory of past inputs that determines the network's current state and response to new information. Similarly, a cell's current state reflects its history of environmental exposures.

Distributed representation occurs because information is not stored in a single location but instead is spread across the entire reservoir, often in a partially redundant manner. This distributed representation provides robustness to noise and damage, a valuable feature in biological networks.

The reservoir network is **tolerant to imperfect wiring** because the random connectivity and stasis mean that precise wiring is not a prerequisite for complex computation.

Simplified learning can happen because RC primarily involves adjusting how the reservoir's activity is interpreted rather than changing the connections within the reservoir itself. In particular, a system uses the stored information in the reservoir to adjust its response or predict a future input. The system can learn or evolve those responses, while the reservoir network itself does not change.

RC typically has emergent functional design at a higher level relative to the partially random nature of lower level components. This architecture means that it may sometimes be difficult to infer function by starting at the lowest level and tracing upwards in a system.

Overlapping circuits may be common because the information stored in an RC network can be read out in different ways for different functions. An existing circuit may provide the basis for new functions by evolving novel responses that read out the circuit's stored information in a different way.

In summary, RC provides an interesting machine learning model for how evolution might have leveraged simple yet powerful design principles to create the complex information processing capabilities observed in living systems.

Reservoirs illustrate five important network attributes, described in more detail below. First, reservoirs have feedback connections among nodes, a kind of recurrent neural network. Recurrence provides the capacity for memory, which allows calculations over past input sequences (Goodfellow et al., 2016). Second, a reservoir's random connectivity transforms inputs to a lower dimensional internal representation.

Third, the circuit often functions precisely in spite of the internal randomness in design (Frank, 2023b).

Fourth, at a small scale, the individual nodes and their connectivity may be nearly uncorrelated with large-scale circuit function.

Fifth, large networks may be overwired. They often have more parameters and connectivity than required for the problem being solved (Frank, 2017; Sorrells & Johnson, 2015).

Machine learning's reservoir model provides analogies for biological circuits. Before turning to those analogies for biology, I summarize additional aspects of reservoir computing.

How do random networks store information?

Think of a liquid with a smooth surface. Drop a stone on the surface to generate waves. Then, drop another stone, and another. The pattern of surface waves contains information about the temporal history of inputs.

A reservoir's computational network acts similarly. Each input propagates signals through the recurrent network based on the random connectivity pattern between nodes and the rules for updating nodes. At each point in time, the network state contains information about the history of inputs (Maass et al., 2002).

Prediction and response

In a computational reservoir model, the circuit can be trained to predict the future values of environmental variables. The randomly connected network takes in a sequence of inputs. The internal states of the network's nodes reflect the past input sequence. The machine learning algorithm then fits a regression model, using the reservoir's current state at any point in time to predict the future values of the environmental variables (Chattopadhyay et al., 2020; Chepuri et al., 2024; Frank, 2023b).

Figure 2 illustrates the response of a reservoir network. The blue curve in (A) traces an input signal. The gold curve shows the network's prediction for input values. Panel (B) magnifies part of panel (A).

Panel (C) clarifies the nature of predictions. At each point in time, the blue curve marks the input value. The gold curve shows the predicted input 2 time units into the future. Panels (A) and (B) shift the gold curve 2 time units to the right so that deviations between the curves measure the difference between the actual and predicted future inputs. Time units are nondimensional and may be interpreted on whatever scale makes sense for a particular application.

Here are the details (Martinuzzi et al., 2022). A simple reservoir network has N internal nodes, with nodal values \mathbf{x}_t at time t. A random matrix, W, describes the connection weights from each node to all nodes. The m input values at each time are \mathbf{u}_t . The matrix \mathbf{W}_{in} describes the connection weight from each input to each network node. The update rule for nodal values is

$$\mathbf{x}_{t+\Delta t} = (1-\alpha)\mathbf{x}_t + \alpha \tanh(\mathbf{W}\mathbf{x}_t + \mathbf{W}_{\mathrm{in}}\mathbf{u}_{t+\Delta t}). \tag{1}$$

The function tanh returns the hyperbolic tangent of its argument. The parameter α sets the rate at which nodal values update. In reservoir models, all parameters remain constant. In later models based on this equation, the learning phase adjusts these parameters.



Figure 1. Reservoir networks provide information about past inputs. The simple examples here illustrate the process. The intensity of shading for each node reflects its value, with lower values having darker shading and darker color. (A) The five nodes in this example create a circular network (Rodan & Tino, 2010). In each timestep, each node passes its value, *v*, transformed by tanh(v/2), to the next node, in which tanh is the hyperbolic tangent function. That function maps its input to the range [-1, 1], keeping the values in the network bounded to a useful set. The node at the top provides the external input into the network, with input value transformed by tanh. In this example, the inputs (0.1,0.3,0.5,0.7,0.9) rise linearly over time. As the inputs proceed from left to right, the network values encode information about the input sequence. The final state on the right has declining values around the circle when starting from the top node, associated with the rising inputs over time. (B) Network dynamics for flat inputs, each of 0.7. In the final state, nodal values decline more slowly. (C) Network dynamics for falling inputs, which are the reverse order of the inputs in (A). In the final state, nodal values are relatively flat, balancing the decline of input values against the decline in values transmitted around the circle. (D) A reservoir network with 10 bidirectionally connected each input to each node are also random. The same network connectivity is used for each instance. The input sufferent for an instance is described at its center. In this case, each input sequence has 11 values. The illustrations show the final state after all inputs. The conclusion is that a network's final state reflects its input pattern, providing the system with information about the history of inputs.

For simple reservoir computing, one fits a regression model that maps the nodal values at a particular time to a future input value. Training a network to make predictions follows standard machine learning protocol. Split a given sequence of inputs into a training part and a test part. Use the training part to fit the model parameters. Use the test part to analyze how well the system performs on data not used for training. See the figure caption for more detail about this particular case.

This example shows that a simple randomly constructed reservoir can store sufficient information about past inputs to make good predictions about future inputs. Reservoir networks have solved a wide variety of challenging problems (Chattopadhyay et al., 2020; Chepuri et al., 2024; Cucchi et al., 2022; Frank, 2023b). For example, such networks predict macroeconomic variables (Ballarin et al., 2024) and spatiotemporal patterns of sea surface temperatures (Walleshauser & Bollt, 2022) and solar irradiance (Li et al., 2020). Other studies suggest that variants of reservoir computing networks provide good models for neurobiological circuits (Loeffler et al., 2020, 2021).

Biological insights from reservoirs

Recurrence, memory, state, and dimension

Reservoirs have four attributes that make them particularly good at using past inputs to predict future inputs. Those four



Figure 2. Match between actual (blue) environmental inputs and predicted (gold) inputs generated by a reservoir network. See text for description of each plot. Briefly, (A) the match over 100 nondimensional time units, (B) magnification of a short time sequence, and (C) gold curve shifted 2 units to the left. I created the reservoir networks with the Julia programming language package ReservoirComputing.jl (Martinuzzi et al., 2022). Each network has 20 nodes. The connectivity matrix was randomly generated with a 0.6 expected frequency of zeros. The connectivity weights were normalized to a matrix spectral radius of 1.0, which means that the nodal values tend neither to increase nor decrease over time. The input is $u(t) = \sin(t) + \sin(0.51t) + \sin(0.22t)$. The total input sequence occurred over time units t = [1,300] in increments of $\Delta t = 0.05$. In the training period during the first 200 time units, I fit a regression model on the nodal values, **x** in eqn 1, to predict the input values at 2 time units into the future, which is $2/\Delta t = 40$ sequence increments. To fit the prediction model, I used lasso regression with an L1 regularization cost of 0.0003. That linear cost on the magnitude of the parameters prevents large parameter values and reduces to zero those parameters that contribute little to the predicted fit. In this case, the fitted regression used 7 of the 20 nodal values, a significant reduction in dimensionality of the input complexity. Reducing the L1 cost to zero used all 20 nodal values in the fitted regression and gave a nearly perfect fit. The plots show the results for the inputs that were not used during the fitting procedure. The freely available Julia computer code provides full details about assumptions and methods for all figures in this article (Frank, 2024b).

attributes are likely to shape the architecture of biological circuits.

Recurrence means that information flows through feedback loops within a network. In machine learning, one thinks of a network as composed of connected nodes, each node contributing calculations to the network's overall computational path.

In the simplest feed forward architecture, information flows unidirectionally, starting from the external inputs and flowing into the first layer of computational nodes. The outputs of that first layer feed forward as inputs to the next layer, and so on, until the final nodal outputs form the circuit's response.

If a node feeds its output back as input to itself, that is a recurrent feedback loop. Such a node gets inputs from different time traces of information flow. A feed forward input represents a relatively current time trace. The self-feedback represents a prior time trace. In general, recurrence occurs when a node receives input from multiple time traces, with feedback loops bringing in earlier time traces.

Recurrence creates a kind of *memory*. With a recurrent loop, the input flowing into a particular node may include both the first trace of information from the most recent external input into the network and traces of information from earlier inputs into the network. The temporal offset created by recurrent flow allows a node to compute on sequences of inputs.

At the network level, memory can arise simply by recurrent connectivity. Each computational node takes inputs and produces outputs without any local memory. Alternatively, nodes may retain memory through their internal *state*. For example, a node that retains memory of its most recent output can use that internal state as input for its next computation. Nodes in reservoir networks may retain internal state, as in eqn 1.

Thus, we have two distinct types of memory. Recurrent network memory is a global property of the network, arising from feedback connections that create the possibility for the inputs into a node ultimately to trace back to external environmental inputs that came into the network at different times. In contrast, local state memory is an internal property of a single node by which the node retains information about past inputs into that node.

Reservoirs perform *dimensional reduction* of their inputs. The temporal input sequence may be a high-dimensional vector. The reservoir's internal nodal states reduce that highdimensional input to a lower dimension.

Because computational reservoirs are typically created with random, unchanging connectivity, many of the internal states may either be redundant or not useful with regard to the particular computational challenge. In reservoir computing, the overall computation picks out the useful states to make a prediction about the challenge, often using regression on the states to predict an outcome. The prediction stage typically reduces the internal state dimensionality to a much lower dimension of useful information. Recurrence, memory, state, and dimension provide a basis for evaluating biological network architecture.

Form and function, discovery, and refinement

Computational reservoirs raise three interesting questions about biology.

First, how does a circuit's form relate to its function? The fact that simple random reservoirs perform well suggests that, for circuits, the low-level connectivity patterns that determine form may often be weakly correlated with function. Good performance arises from the nonspecific attributes of recurrence, memory, state, and dimensional reduction, which transcend the specific details of connectivity between nodes.

Second, how do things get started? Natural selection can only refine the performance of something that already exists and has the capacity for minor adjustments to improve function. Random reservoirs may be a first step in creating functional circuits, providing the basis for something that initially works and can subsequently be improved.

Third, if initial discovery arises by a network that is randomly connected, how effective is natural selection at subsequently refining the circuit architecture? It may be that random connectivity happens only in the initial phases of a circuit's evolution. Subsequently, selection may refine the circuit, hiding the initial randomness from view. Or the initial random connectivity may be sufficiently good, making further selection too weak to modify form.

Overall, how often do circuits first arise by co-opting an existing network that is initially random in its connectivity with respect to function? How often will subsequent selection erase traces of that initial randomness? In other words, what sorts of relations between form and function should we expect at different stages in a circuit's evolution? The following examples and comments suggest how we might start on these questions.

Origin of traits: perception and response

The origin of perception-response traits poses a puzzle. Perception provides no benefit without a response. A response cannot happen without perception. How do things get started? That question guides how we may think about form and function, discovery, and the refinement of biological circuits.

Reservoir computing suggests a solution. If, initially, a randomly connected network provides information about environmental inputs, then a biological circuit can subsequently use that perceived information to evolve an appropriate response (Frank, 2023b). Once a working perception-response circuit is in place, natural selection may refine the circuit.

The prior section showed how a random reservoir might initiate a new perception-response circuit. In Fig. 2, the system takes in a single input that fluctuates in a complexly periodic way over time. A sparsely and randomly connected network with 20 nodes stores an imperfect and dimensionally reduced memory of past inputs. That memory provides a sufficient basis for a population of organisms to learn, by natural selection, how to predict future environmental states using only the internal states of its randomly connected network. Many studies show that random networks provide the foundation for predicting complex patterns of future environmental fluctuations (Chattopadhyay et al., 2020; Chepuri et al., 2024; Cucchi et al., 2022; Frank, 2023b; Loeffler et al., 2020, 2021). Given that random networks often perform very well, how much might we expect natural selection to refine such circuits? There is a tendency in the systems biology literature to expect that natural selection refines every small-scale feature of a biological circuit. However, at a small scale, the intensity of natural selection may often be rather weak relative to various stochastic evolutionary processes. Circuits may often be more randomly or more complexly wired than expected from a purely optimizing learning process (Frank, 2023c; Lynch, 2024; Muñoz-Gómez et al., 2021).

Prokaryotes, eukaryotes, and multicellularity

When a novel challenge arises, organisms may benefit from a pre-existing network that is randomly connected with respect to the new challenge. What might influence the chance that there is a usable pre-existing network?

Comparing prokaryotes and eukaryotes provides an interesting contrast. Prokaryotes typically have smaller cells that are more limited in their total protein content. Proteome limitation may constrain the transcription factor network that comprises the primary cellular response circuits (Chen et al., 2021; Frank, 2022a). In addition, populations of prokaryotes may typically be larger than populations of eukaryotes. Larger populations mean that natural selection is more likely to outweigh the random sampling processes of small populations, more effectively refining the details of functional circuits and reducing the amount of overconnected random wiring (Frank, 2023c; Lynch, 2007).

Prokaryotic transcription factor networks may therefore be less widely connected and more finely pruned than eukaryotic networks. If so, then prokaryotes may be less likely to have pre-existing networks that are essentially random with respect to new challenges. And prokaryotes may be less likely than eukaryotes to retain random or excessively wired networks after the refining action of natural selection.

Comparing multicellular and unicellular eukaryotes, two factors favor broader and more randomly connected networks in multicellular organisms. First, multicellular populations tend to be smaller, with more random fluctuations that weaken the relative efficacy of natural selection (Lynch, 2007). Second, in multicellular organisms, individual cells are tuned for the differentiation of functions rather than the unicellular tuning for fast growth and efficient resource usage. The greater emphasis on speed and efficiency in unicells likely prunes the size and randomness of their circuits.

Overall, broader circuit connectivity and more randomness likely occur in multicellular eukaryotes, followed by unicellular eukaryotes, and then unicellular prokaryotes. Perhaps the diversity and complexity of multicellular eukaryotic cellular circuits provides greater opportunity for the discovery of novel functions.

Precise traits from sloppy components

A trait may first arise from a circuit with reservoir-like random connectivity. Subsequently, natural selection may not be sufficiently strong to refine and prune the network. A circuit with reservoir-like random connectivity appears rather sloppy in its low-level design yet can function precisely with respect to its environmental challenge.

In other words, we may see precise traits arising from sloppy components (Frank, 2023c). Other evolutionary pathways can also lead to low-level sloppiness of circuits that nonetheless function in a relatively precise way (Daniels et al., 2008; Frank, 2007, 2023c; Gray et al., 2010; Gutenkunst et al., 2007; Lynch, 2012, 2024).

Network architecture

Random connectivity can be effective for certain challenges. However, within reservoir computing, imposing some structure to connectivity patterns often improves circuit performance or reduces the size of the required reservoir (Loeffler et al., 2020, 2021; Rodan & Tino, 2010; Rodan & Tiňo, 2012).

In biological circuits, hierarchy and modularity may occur within pre-existing circuits. Or such broad-scale structure may evolve to refine circuit performance (Tripathi et al., 2023). For example, *Caenorhabditis elegans* organizes its 302 neurons into a modular architecture, which may enhance the computational efficacy of its neural circuits (Ruach et al., 2023).

Neurobiology

Machine learning models have long been linked to neurobiological circuits (Hebb, 1949; Rumelhart & McClelland, 1986; McCulloch & Pitts, 1943). Recent articles continue to develop those connections (Jiahui et al., 2023; Kanwisher et al., 2023), including discussion of reservoir computing as a model for particular aspects of neurobiology (Ramezanian-Panahi et al., 2022). Although this article emphasizes genetic and biochemical circuits of cells, it is helpful to mention a few applications of reservoir computing concepts to neural circuits.

For small neural circuits, the connectivity of reservoir wiring may be relatively repeatable from one individual to another. With developmentally predictable wiring, the readout of the reservoir states can also evolve to be fixed developmentally. For larger circuits, the reservoir wiring may vary developmentally between individuals. The internal reservoir states no longer express predictable information about inputs to the circuit.

In the case of larger circuits with less consistent wiring, the organism must learn to read the circuit's states in relation to the particular environmental problem that the circuit must solve. Learning to read a reservoir circuit may be the origin of commonly observed critical learning periods (Frank, 2023b).

At a more abstract level, Czégel et al. (2021) ask how complex intelligence may arise from underlying neural circuits. They studied a computational model of initially random reservoir circuits within a single brain that, in effect, teach one another. The ensemble of circuits uses Darwinian trial-and-error selection to learn the solutions to complex challenges.

Dimensional reduction

Systems often encode high-dimensional inputs to a lowdimensional representation, highlighting important attributes (Reddy et al., 2020; Sorzano et al., 2014). Systems may also contain a model of the environment, which helps to transform an internal representation of inputs into an appropriate action (Bin et al., 2022; Gupta & Khammash, 2023; Schölkopf et al., 2021). Box 2 provides general background for this topic.

Box 2: Dimensional reduction

Understanding how cells process complex environmental information into regulatory responses presents a fundamental challenge in biology. Machine learning, particularly the concept of dimensional reduction, offers powerful analogies for understanding how cellular networks accomplish this feat. Just as machine learning systems must distill vast amounts of complex data into meaningful patterns, cells must convert numerous environmental signals into specific regulatory decisions.

Dimensional reduction refers to the process of converting high-dimensional data into simpler, lower-dimensional representations while preserving essential information. Think of how a skilled artist can capture the essence of a complex scene in just a few brushstrokes. The artist has learned which features are essential to convey meaning and which can be omitted. Similarly, both machine learning systems and cells must learn to extract and represent the most important aspects of their environment while filtering out noise and redundancy.

In machine learning, encoders provide an elegant solution to this challenge. An encoder is like a skilled translator who must convey a story in fewer words while maintaining its essential meaning. The encoder learns to compress information by identifying patterns and relationships in the input data, creating a compact code that captures the key features. The power of this approach lies in how encoders discover efficient representations automatically through experience. They learn which combinations of features matter most for predicting outcomes.

For example, an encoder processing images of faces might learn to track key features such as the distance between the eyes or the shape of the jaw, rather than storing every pixel. This compact representation contains the essential information needed to recognize or reconstruct the face. Similarly, cellular signaling networks may act as biological encoders, learning over evolutionary time to track key combinations of environmental signals rather than responding to each input independently.

The architecture of artificial encoders may provide clues for how cellular signaling cascades compress environmental information into key molecular signals, which regulatory networks then expand into broader cellular responses. Just as machine learning encoders must balance the trade-off between compression and preservation of important information, cellular networks face similar constraints in processing environmental signals efficiently with limited molecular resources.

In both machine learning and cellular systems, dimensional reduction through encoding offers several key advantages. It filters out noise and irrelevant variations, **making responses more robust**. It reduces the complexity of information processing, making it more energetically efficient. Perhaps most importantly, it **allows systems to generalize** from past experiences to new situations. Just as a machine learning model can recognize patterns in never-beforeseen data, cellular networks can mount appropriate responses to novel combinations of environmental signals.

Viewing cellular regulatory networks through the lens of dimensional reduction provides valuable insights for biological research. It suggests new ways to analyze and model cellular decision-making, helps predict cellular responses to complex stimuli, and offers frameworks for understanding how cells maintain robust responses despite noisy inputs. This perspective also helps explain how relatively simple molecular networks can process complex environmental information and generate appropriate cellular responses.

Encoder network

An encoder network in machine learning typically processes a high-dimensional input within its initial layer. That first layer then passes a lower-dimensional output to the next layer of the network. A series of dimension-reducing layers may follow, each layer producing fewer outputs than the number of inputs into the layer (Chen & Guo, 2023; Mao et al., 2024).

By forcing the system to reduce dimensionality through a series of partial steps, each reducing layer has the opportunity to capture different essential features of the initial input. The sequential reduction sometimes finds efficient and effective representations that can be difficult to find when reduction is attempted in a single step.

The input may occur as a temporal sequence. If so, then the initial network layer stores a memory of past inputs in order to find an effective reduction of the input sequence. A memory trace can be created by recurrent network connections, like those described by eqn 1.

Predicting direction of change

Suppose, for example, that a system experiences a sequence of environmental inputs. The system gains by predicting the direction of change for the next input relative to the current environment. Prior inputs provide information on the next direction of change.

I develop an example for which we can obtain the optimal dimensional reduction and internal decision model. We can then observe how a machine learning model discovers and encodes a nearly optimal solution.

For many realistic problems, we may not be able to solve for or guess what sort of circuit solution is required. In those cases, machine learning guides us in the design of circuit solutions and the analysis of how biological circuits solve environmental challenges.

Looking out toward those more difficult problems, the basic example developed here illustrates three points. First, very simple networks can learn to encode efficient computational circuits based solely on feedback about current performance. Second, dimensional reduction is a common attribute of effective circuits. Third, circuits may encode a causal model rather than exploit a database of correlations for past inputs. This example extends a previous study (Frank, 2024a).

Optimal solution

To generate input sequences, start with a random walk, \hat{u}_t . Then replace each value at time *t* by its exponential moving average, $u_t = \beta \hat{u}_t + (1 - \beta)u_{t-1}$. Figure 3A and E show two example sequences. The figure legend describes details.

The widespread biological applicability of this process can be seen by writing an approximately equivalent differential equation, $\dot{u} = \hat{u} - \lambda u$, in which \dot{u} is the change with respect to time for some biological quantity, \hat{u} is an input that drives change in this quantity, and λ is the intrinsic decay rate of u. For this differential equation, the solution for u at time t is proportional to the continuous time exponential moving average of u at that time.

Here, we return to the discrete form, which is easier to work with for our purposes. The goal is to predict the direction of change, the sign of $\Delta u_t = u_{t+1} - u_t$. The sign is positive when

$$\Delta u_t = \beta \Delta \hat{u}_t + (1 - \beta) \Delta u_{t-1} > 0$$

Because \hat{u}_t is a random walk, its change $\Delta \hat{u}_t$ is equally likely to be positive or negative. Thus, the best prediction for the sign of Δu_t is the sign of Δu_{t-1} . In other words, the most recently observed direction of change provides the best prediction for the next direction of change.

Roughly speaking, we may think of the recently observed change, Δu_{t-1} , as the trend momentum. Positive momentum means that the trend is likely to continue up. Negative momentum means that the trend is likely to continue down. Later comments about momentum refer to this term.

The greater the momentum, the more likely the trend will continue in the same direction. For example, the greater Δu_{t-1} , the more negative the underlying random walk change, $\Delta \hat{u}_{t}$, must be to reverse the trend. Increasingly extreme moves in the underlying random walk are increasingly uncommon.

Thus, an ideal internal model uses the currently observed trend direction to predict the next direction of change. And it uses the trend momentum to estimate the confidence in the directional prediction.

Circuit that predicts directional change

A computational circuit can learn to solve this directional prediction problem. The initial layer of the circuit has the recurrent memory structure in eqn 1 with a pool of $N = 2^n$ internal states.

For $n \ge 2$, I used a sequence of layers with $2^n, 2^{n-1}, ...$ for dimensional reduction. In some cases, larger networks learned more quickly. But an initial layer with two internal memory states was sufficient. The circuit needs information only for the current and prior inputs to generate an estimate for the momentum, Δu_{t-1} .

I focused on this smallest two-state circuit, as shown in Fig. 4. The figure legend describes how the circuit transforms inputs into two internal memory states, reduces those memory states to an internal one-dimensional summary value, and then transforms the reduced summary value into a prediction about future change.

Figure 3 shows results for one particular circuit. That circuit learned to predict the direction of change by adjusting its 12 internal parameters. The adjustment happened by a standard machine learning cycle. In each cycle, I generated a sample input such as the one shown in Fig. 3A. I then calculated the



Figure 3. Encoder circuit prediction for the direction of change in a sequence of observations. The input sequence is calculated by starting with a random walk, $d\tilde{u} = \sigma dW$, at $\tilde{u}_0 = 0$ with $\sigma = 0.2$, in which dW is a Wiener process that samples a normal random variable with standard deviation of \sqrt{dt} , with dt = 0.01 in all examples. The sequence values are normalized to [0,1] by affine transformation, yielding \hat{u} . The sequence is then replaced by its exponential moving average, $u_t = \beta \hat{u}_t + (1-\beta)u_{t-\Delta t}$, sampled at discrete time points, and with $u_0 = \hat{u}_0$ and $\beta = 0.2$. Figure 4 shows the circuit architecture. The 12 network parameters in the Fig. 4 circuit were adjusted to reduce the loss function by the Adam learning algorithm with learning rate 0.005 applied to 25,000 randomly generated sequences. The examples in this section analyze 300 time units with 10 sample points per time unit for a total of 3,000 sample points. After optimization, new sequences were used to test performance, as follows. (A) Input sequence. (B) The fitness for each prediction is the absolute value of the difference between the current and prior observation multiplied by -1 if the prediction about direction is incorrect. Cumulative fitness is proportional to the sum over all predictions up to the current time point. (C, D) A magnified view of a time interval from the plots above. (E–H) Similar plots for a second input sample.

performance by using the loss function described in the figure legend.

Initially, a circuit predicts the direction of change with a success rate of 0.5, equivalent to flipping a coin. The computer code then calculates a gradient for how changes in the parameters alter the performance. The algorithm adjusts the parameters to improve performance slightly. The process repeats with a new random input sequence. In this example, I used 25,000 rounds of parameter refinement.

Figure 3B and F show the performance of the circuit after the final parameter refinement, measured as cumulative

fitness, as described in the figure legend. An upward trend in cumulative fitness associates with an accuracy of prediction for the direction of future change that is greater than the random baseline of 50%.

The method that generated the random sequences creates inputs for which there is an $\sim 80\%$ match between the direction of change at one time and the direction of change at the next time. Figure 5 shows that the circuit correctly predicted the direction of change at nearly the optimal match frequency.

Figure 6 describes the mechanism that the circuit uses to predict the direction of change. One internal memory state



Figure 4. Encoder dimensional reduction network for predicting sequence trend. In this example, the input is a sequence of values. The recurrent memory layer is described by egn 1, allowing the network to calculate information about current and past inputs and retain that information in the internal states, x_i . Output from this layer is then passed as input to the dimension reduction layer. That layer sums an affine transformation of each input to produce a 1D output passed to the final layer. The final layer applies to its input an affine transformation followed by a sigmoid function, transforming inputs on $(-\infty, \infty)$ to the final output on [0,1]. The output, ρ , is taken as a prediction of the direction of change of the next input value relative to the current input value. An actual positive change in the data associates to a target of $\kappa = 1$, and a negative change associates to a target of $\kappa = 0$. The distance between the prediction, ρ , and its target, κ , is the cross-entropy loss function, $-\kappa \log \rho - (1-\kappa) \log (1-\rho)$. Figure 3 shows an application of this circuit architecture to a particular example.

tracks the prior input. The other internal state is proportional to the momentum. Those states provide sufficient information to generate nearly optimal predictions about direction. The circuit also has information about its relative confidence in its predictions, given by the magnitude of the momentum estimate.

Internal model

Consider a trending numerical sequence, as in the previous examples. A circuit could learn to predict the direction of change in the next observation in two different ways.

First, the circuit could exploit patterns of correlation. If, for example, in the data from which the circuit learns, 8 typically follows the target sequence 2–4–5–9, then the circuit may learn to predict a negative change in response to the target.

learned correlation 2 4 5 9
$$8 \implies -$$

Second, the circuit could learn an internal causal generative model for the sequence. For example, instead of focusing on the correlation between particular trigger sequences and the direction of future change, the circuit might learn the sort of causal model described in prior sections.

causal model 2 4 5 9
$$X \implies +$$

Here, *X* is predicted to be larger than 9 because of the general tendency for trends to continue in their current direction. If a very simple circuit with limited memory can learn this underlying causal model, then it can predict the direction of change with nearly optimal accuracy.

Learning algorithms easily exploit correlations. If those correlations are just sampling artifacts in the original data, then the circuit will perform poorly in new situations. Sampling artifacts happen frequently in datasets that are small relative to their dimensionality, sometimes causing circuits to be fragile when focused on correlations.

Alternatively, if the learning dynamics can successfully reduce the dimensionality of the data in relation to the circuit's goal, then the circuit may be able to discover a relatively simple model that follows the contours of the underlying causal process. A good internal causal model performs well in response to novel input data.

We do not know how often simple cellular circuits encode internal models of the environment rather than correlations between environmental states. An interesting literature develops the claim that internal models are particularly powerful solutions for prediction and control (Bin et al., 2022; Gupta & Khammash, 2023). Of course, given complete data on all possibles sequences and perfect predictions by correlational estimates, a correlational network converges to a perfect internal model. But, in practice, small computational networks may be able to calculate outcomes simply and quickly, whereas learning and computing highly accurate correlational predictions may be very costly or impossible.

A biochemical circuit for trend prediction

This paper argues that machine learning and biological evolutionary learning may solve similar challenges with similar kinds of circuits. For example, the previous subsections showed how machine learning uses dimensional reduction to solve the problem of trend prediction. In this subsection, I use the insights from the machine learning solution to develop a hypothesis for how a simple biochemical circuit might solve the same problem.

In the machine learning solution, an internal model reduces the high dimensional input sequence to a single difference between the internal state estimates for the current input and a recent input. The sign of that difference predicts the direction of future change. The magnitude of that difference estimates the confidence in the prediction.

To build a biochemical circuit, we start with the machine learning solution's calculation of the difference between internal estimates for inputs at different time points. The following system of differential equations expresses a candidate biochemical circuit

$$\dot{x} = \alpha u - \beta x$$

$$\dot{y} = \gamma (\alpha u - \beta y)$$

$$\dot{z} = \lambda + \eta (y - x) - \delta z.$$
(2)

Overdots denote derivatives with respect to time, t. The variables x, y, z, u are functions of time. Here, u is a randomly changing input that follows the same process as described in the caption of Fig. 3.

For constant input, u, the molecular of abundances of xand y have the same equilibrium values, α/β . As the input, u, changes, *y* changes more slowly than *x* when $\gamma < 1$. Thus, y-xtends to be negative when u is increasing and positive when u is decreasing. Therefore, the deviation of z from its equilibrium, λ/δ , estimates the direction and magnitude of recent changes in *u*, providing a prediction for the direction of future change.

Using the methods for input generation and discrete time point sampling in the caption of Fig. 3, I optimized the six parameters in eqn 2. Figure 7 shows the results for a particular set of optimized parameters and a single realized sequence of the random input variable, *u*.

In this model, the internal states x and y approximately track a shorter and a longer exponential moving average of the inputs. The difference between those moving averages anticipates future trend.



Figure 5. Circuit accuracy in predicting the direction of input change. Random input sequences were generated as described in Fig. 3. For a particular generated sequence with 100,000 sample points, the direction of change between two inputs predicted the direction of change for the next input with frequency 0.795, which estimates the maximum accuracy that could be achieved on that sequence. For each of 10,000 novel sequences with 3,000 sample points, I calculated the deviation between the optimized circuit's frequency of correct predictions and the maximum estimated accuracy. The histogram shows the distribution of those deviations. The median, mean, and standard deviation refer to that distribution of deviations.



Figure 6. Mechanism used by the machine learning encoder circuit to predict the direction of change in inputs. A sequence of inputs was generated as described in Fig. 3. These plots show a subset of the sample points. The circuit has two internal memory states, described in Fig. 4. (A) The blue curve traces the prior input value. The gold curve shows one of the encoder circuit's internal memory states multiplied by a constant. (B) The other internal state rises and falls with the estimated momentum of the trend, which is the difference between the current and prior input, shown in gold. I transformed the state value to have the same mean as the input curve, which associated a positive internal momentum estimate (gold) with a location above the input (blue) curve and a negative momentum estimate with a location below the input curve. The blue curve is the same as in the upper panel. (C) The predicted direction rises and falls with a function of the momentum. The sign provides the predicted direction. The magnitude describes an estimate for the confidence of the prediction. The actual circuit outputs predicted values *z* on the interval (0,1). Here, I show 2(z - 0.5) so that positive and negative values correspond to positive and negative predicted changes.



Figure 7. Biochemical circuit prediction for the direction of change in a sequence of observations. Dynamics described by eqn 2. Parameters optimized by the method described below. (A) Molecular abundance of *x* compared with the input sequence, *u*, in blue and gold, respectively. I multiplied *u* by a constant to show the proportional match in the dynamics. The abundance of *x* provides a good internal estimate of the recent input value. (B) Comparison of *x* and *y*, in blue and gold, respectively. The abundance of *y* responds more slowly to input changes. Thus, y - x estimates the recent difference in input values. (C) The abundance of *z* provides a prediction for the direction of change in the next sampled input. This panel shows the deviation of *z* from its equilibrium multiplied by the parameter ϕ , as $\rho = \phi(z - \lambda/\delta)$. The sign and magnitude of ρ predict the direction of change and the confidence in the prediction. In this example run, the sign of ρ predicted the sign of the next input change with 72.2% accuracy compared with a true match between the sign of the difference in one increment and the sign of the difference in the next increment of 72.7%. I optimized the six parameters in eqn 2 by minimizing the cross-entropy loss described in the caption of Fig. 4. To minimize the loss function, I obtained values for the six parameters of eqn 2 and for ϕ by the Sophia algorithm with learning rate 0.005 applied to 10,000 randomly generated input sequences (Liu et al., 2024). Each input sequence was generated over 10 time units and sampled at 0.1 time units, yielding 101 points. The times between points were interpolated linearly. To apply that input sequence to eqn 2, the time scaling for the inputs was multiplied by 3,000 to yield a time span of 30,000 sampled at intervals of 300 time unit intervals to obtain the 101 sample points. In each run of the minimization routine and in these plots, I skipped the first 10 sample points.

By adjusting the memory window of those moving averages through the parameters β and γ , the system can modulate the timescales of trend estimation, future prediction, and the tradeoff between the overall accuracy and false signal frequency for trend reversals. Thus, this circuit provides a very simple and general way in which organisms can predict future environmental trends based on past environmental inputs (Frank, 2024a). In retrospect, the structure of this anticipatory circuit seems obvious. However, I did not guess the solution until I observed how the machine learning circuit solved this problem.

For more complex problems, following the same method may often be helpful. First, use machine learning to search for how a circuit might solve the problem. Second, analyze how the machine learning circuit works. Does it embed an internal model? Third, construct and optimize a biological circuit that works in the same way. Finally, use the architecture of the biological model as a hypothesis for how actual biological circuits may be designed to solve the problem.

For the relatively small machine learning circuits that make good analogies for biochemical circuits, we can often reverse engineer the computational mechanism, as in this example.

Alternatively, we could start with a general system of differential equations that include a large pool of potential biochemical circuits. We then use machine learning methods directly on the system of equations to search for good biochemical circuit solutions to the specific challenge (Frank, 2022b, 2023a; Hiscock, 2019).

Dimensional reduction in biology

Effective circuits reduce input dimension. Borges made this point clearly when, in his story "Funes the Memorious," the title character recalled every detail of every moment of every day. Funes was incapacitated by his lack of reduction. The story's narrator concluded that "To think is to forget a difference, to generalize, to abstract." O'Gieblyn (2024) summarized Funes incapacity:

Because Funes is incapable of forgetting any detail, he has problems understanding language. He cannot fathom why the word "dog" is used for so many distinct dogs, all of which vary in size, color, and form. He cannot even understand why the same dog is called by a single name why, for instance, "the dog at three fourteen (seen from the side) should have the same name as the dog at three fifteen (seen from the front)."

To learn is to exclude what does not matter. Dimensional reduction shaves to the invariant core. The issues concern: What is the best reduction for a given challenge? How is the reduction process embedded within the circuit? How do circuits learn to accomplish such reduction? Machine learning may become a helpful tool to answer these questions for biology because machine learning and biology face similar challenges (Bengio et al., 2013; Chater & Brown, 2008; Kemp & Tenenbaum, 2008; Levin, 2024; Quiroga et al., 2005; Tishby & Polani, 2011).

This section's examples illustrate the role of dimensional reduction in biological circuits. They also hint at the relations between the biological circuits discovered by natural selection and the engineered circuits discovered by machine learning (Eckmann & Tlusty, 2021; Gamoran et al., 2024; Kanwisher et al., 2023; Mao et al., 2024).

Metabolism, immunity, and development

Bacterial cells switch metabolic pathways based on the internal concentrations of a few key molecules. These molecules serve as an internal representation, condensing the highdimensional state of many nutrients into a simpler form used by the cell (Kochanowski et al., 2017; Wu et al., 2022).

Many biological circuits have an hourglass architecture. Diverse inputs cascade through the initial circuit layers toward a narrow middle processing layer. The diverse inputs may be thought of as the broad base of the hourglass, leading upward toward a narrow central waist. From the limited diversity in the narrow waist, an increasing number of different types flow upward toward the broad output layer at the top (Csete & Doyle, 2004, 2002; Doyle & Csete, 2011; Frank, 2023d; Kaneko, 2024). The narrow waist encodes a dimensionally reduced representation of the inputs, which the circuit decodes into a higher dimensional output. Friedlander et al. (2015) summarize examples of hourglass architectures in metabolism, immunity, and other biological pathways.

Metabolism starts with diverse nutrient inputs. Catabolic processing reduces those inputs to 12 common intermediates, such as pyruvate and fructose-6-phosphate. From those 12 intermediates, a cell makes most of its carbohydrates, nucleic acids, proteins, and other diverse biochemical forms (Csete & Doyle, 2004; Zhao et al., 2006).

Innate immune responses can be triggered by a wide diversity of molecular patterns associated with invasion. The immune system reduces those inputs to changes in a small number of toll-like receptors, interleukins, and a few other molecular types. That reduced immune encoding controls the broad response in gene expression and effector molecules that defend against attack (Beutler, 2004; Oda & Kitano, 2006; Tang et al., 2012).

Development often flows through an hourglass circuit (Frank, 2023d; Irie & Kuratani, 2014; Jordan & Miska, 2023; Shook & Keller, 2008). Early developmental stages vary widely, forming the broad base of the hourglass. Intermediate stages reduce to fewer dimensions, the constrained waist. Later stages diversify to broad variations in adult form.

In the nematode *C. elegans*, the developmental hourglass constrains how proteins explore the evolutionary landscape (Ma & Zheng, 2023). Those proteins that affect early development evolve relatively rapidly when compared with related nematode species, the diverse hourglass base. Proteins that affect intermediate stages evolve relatively slowly, the dimensionally reduced hourglass waist. Proteins that affect late stages evolve relatively rapidly, the diverse hourglass top.

Doyle argued that this widely observed hourglass architecture arises inevitably in all complex, robust circuits (Csete & Doyle, 2004, 2002; Doyle & Csete, 2011; Frank, 2023d). The dimensionally reduced narrow waist defines protocols that connect different layers in the processing flow.

For example, in metabolism, many different biochemical inputs are broken down into the same limited set of protocollike core molecules. Those core molecules provide the basis for building new molecular types. The acquisition and breakdown of the building blocks create one layer, separate from the second layer that builds up functional forms. The digestion layer and buildup layer are connected by a dimensionally reduced common protocol core. The universality of metabolism arises from that reduced core.

Neurobiology circuits

We have the complete neural wiring diagram of *C. elegans*. That connectome has an hourglass circuit architecture (Batta et al., 2021; Sabrin et al., 2020). Information from about 90 sensory neurons flows into a middle layer of ~80 recurrently connected interneurons. The third layer of about 120 motor neurons provides much of the network's response to inputs.

The middle layer includes 10–15 highly connected neurons that form the core hourglass waist, in which sensory inputs compress to a low dimensional encoding. Action follows from decoding the compressed information in the circuit waist, linking the broad sensory input and functional output layers through the narrow core protocol layer.

Rodent neural networks, vastly greater in size and connectivity, also have an hourglass circuit architecture (Stuber, 2023). The sensory cortex takes in high dimensional information directly from all senses and from various intermediate processing ganglia. Those inputs are then compressed sequentially through a series of layers, including the limbic and frontal cortices, the hypothalamus, and the extended amygdala, leading finally to the greatest dimensional reduction in the core ventral tegmental area (VTA) dopamine layer.

The core VTA dopamine layer encodes the primary motivational protocols that transform the initial sensory inflow into a dopamine encoding for goal-directed behavior. The path from the core dopamine protocols to behavior flows through several decoding layers that expand dimensionality, including the ventral striatum, basal ganglia, and corticospinal motor circuits.

In machine learning, multilayer encoding and decoding subsystems are often beneficial for high-dimensional challenges. Each layer separately reduces or builds particular key dimensions in the encoder–decoder process that are difficult to learn or process simultaneously within single layers (Goodfellow et al., 2016).

A study of human facial recognition compared how computational and biological networks parse facial features and differentiate individual faces (Jiahui et al., 2023). Both networks converged to essentially the same dimensional reductions for facial representation geometries and categorical face attributes. However, the networks differed in how they mapped the reduced encoding of facial features to the identification of individual faces.

The authors concluded that computational networks "provide a good model for early cognitive and neural face processing of categorical attributes but are a poor model for individuation."

I would rephrase that conclusion. A programmer makes many specific choices when building a computational model. The authors found that their particular computational choices for a deep convolutional neural network provided a good match for the neurobiological processing of facial categorical attributes. However, their particular computational implementation separated faces differently from how the neurobiological network separated faces.

It is likely that different implementations of computational models would differ in the degree to which they match the biological networks in categorical attributes and individuation. We might learn a lot about how both computational and biological networks function by studying how various computational implementations differ in their match to biology. It often happens that initial failures ultimately reveal the greatest insight.

For example, the ultimate performance metric in biology likely differs from the performance metric that a programmer might initially choose to make a computational model. Specifically, biological design ultimately depends on reproductive success, whereas a programmer would likely use some measure for the success in parsing individual facial components or the success in discriminating between individuals. Reproductive success probably does not map linearly to a programmer's measures of facial parsing and discrimination.

It would be interesting to know how different performance metrics alter machine learning circuits. That type of study might help to understand how different aspects of dimensional reduction in facial or other input data relate to function, providing testable hypotheses about biological circuits.

Going forward

We know how some biological circuits work. But we know relatively little about the full range of life's circuits. Reverse engineering is hard.

For some circuits in prokaryotes, we can list many of the molecules involved and their interaction partners. But the actual circuit designs with respect to information processing, dimensional reduction, and logic at a functional level remain open problems for many prokaryotic cellular processes. Eukaryotic cells and tissues are more complex.

Sometimes, it is useful to take stock of what we do not know. Rapid advances in machine learning provide a basis for inventory, opening a way to explore how various circuit designs handle different challenges. In addition, we can easily alter constraints and performance metrics in machine learning to see how those factors change favored circuits. Such methods provide powerful tools for building a conceptual foundation and generating novel hypotheses.

Differences between machine learning and biology also provide interesting conceptual challenges. Are there intrinsic differences between natural selection and other learning algorithms in how they discover and design circuits? How do differing goals alter architecture?

Synthetic biology sets another challenge (Braccini et al., 2023). How do we design biochemical circuits that achieve particular functions and performance goals? Machine learning provides a powerful set of tools. To use those tools effectively, we must figure out how to translate machine learning circuits into biochemical circuits or how to use machine learning tools directly to design synthetic biological networks (Frank, 2022b, 2023a; Hiscock, 2019; Mishra et al., 2023; Nielsen et al., 2016; Wu & Tan, 2023).

Data availability

Software to produce the figures is available at https://github. com/evolbio/Circuits_01.

Author contributions

S.A.F. did everything.

Funding

The Donald Bren Foundation, US Department of Defense grant W911NF2010227, and US National Science Foundation grant DEB-2325755 support my research.

Acknowledgments

For Boxes 1 and 2, I used various large language model systems to edit or generate parts of the text, including Chat-GPT 40, ChatGPT 01-preview, Claude 3.5 Sonnet (New), and Gemini Pro 1.5.

References

Aldarondo, D., Merel, J., Marshall, J. D., Hasenclever, L., Klibaite, U., Gellis, A., Tassa, Y., Wayne, G., Botvinick, M., & Ölveczky, B. P. (2024). A virtual rodent predicts the structure of neural activity across behaviours. *Nature*, 632(8025), 594–602.

- Ballarin, G., Dellaportas, P., Grigoryeva, L., Hirt, M., van Huellen, S., & Ortega, J.-P. (2024). Reservoir computing for macroeconomic forecasting with mixed-frequency data. *International Journal of Forecasting*, 40(3), 1206–1237.
- Baltussen, M. G., de Jong, T. J., Duez, Q., Robinson, W. E., & Huck, W. T. S. (2024). Chemical reservoir computation in a self-organizing reaction network. *Nature*, 631(8021), 549–555.
- Batta, I., Yao, Q., Sabrin, K. M., & Dovrolis, C. (2021). A weighted network analysis framework for the hourglass effect—and its application in the C. elegans connectome. PLOS ONE, 16(10), e0249846.
- Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8), 1798–1828.
- Beutler, B. (2004). Inferences, questions and possibilities in toll-like receptor signalling. *Nature*, 430(6996), 257–263.
- Bin, M., Huang, J., Isidori, A., Marconi, L., Mischiati, M., & Sontag, E. (2022). Internal models in control, bioengineering, and neuroscience. *Annual Review of Control, Robotics, and Autonomous Systems*, 5, 55–79.
- Braccini, M., Collinson, E., Roli, A., Fellermann, H., & Stano, P. (2023). Recurrent neural networks in synthetic cells: a route to autonomous molecular agents? *Frontiers in Bioengineering and Biotechnology*, 11, 1210334.
- Chater, N., & Brown, G. D. (2008). From universal laws of cognition to specific cognitive models. *Cognitive Science*, 32(1), 36–67.
- Chattopadhyay, A., Hassanzadeh, P., & Subramanian, D. (2020). Datadriven predictions of a multiscale lorenz 96 chaotic system using machine-learning methods: Reservoir computing, artificial neural network, and long short-term memory network. *Nonlinear Processes in Geophysics*, 27(3), 373–389.
- Chen, S., & Guo, W. (2023). Auto-encoders in deep learning—a review with new perspectives. *Mathematics*, 11(8), 1777.
- Chen, Y., van Pelt-KleinJan, E., van Olst, B., Douwenga, S., Boeren, S., Bachmann, H., Molenaar, D., Nielsen, J., & Teusink, B. (2021). Proteome constraints reveal targets for improving microbial fitness in nutrient-rich environments. *Molecular Systems Biology*, 17(4), e10093.
- Chepuri, R., Amzalag, D., Antonsen, T. M., & Girvan, M. (2024). Hybridizing traditional and next-generation reservoir computing to accurately and efficiently forecast dynamical systems. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 34(6), 063114.
- Csete, M., & Doyle, J. (2004). Bow ties, metabolism and disease. Trends in Biotechnology, 22, 446–450.
- Csete, M. E., & Doyle, J. C. (2002). Reverse engineering of biological complexity. *Science*, 295(5560), 1664–1669.
- Cucchi, M., Abreu, S., Ciccone, G., Brunner, D., & Kleemann, H. (2022). Hands-on reservoir computing: A tutorial for practical implementation. *Neuromorphic Computing and Engineering*, 2, 032002.
- Czégel, D., Giaffar, H., Csillag, M., Futó, B., & Szathmáry, E. (2021). Novelty and imitation within the brain: A darwinian neurodynamic approach to combinatorial problems. *Scientific Reports*, 11(1), 12513.
- Damicelli, F., Hilgetag, C. C., & Goulas, A. (2022). Brain connectivity meets reservoir computing. *PLoS Computational Biology*, 18(11), e1010639.
- Daniels, B. C., Chen, Y.-J., Sethna, J. P., Gutenkunst, R. N., & Myers, C. R. (2008). Sloppiness, robustness, and evolvability in systems biology. *Current Opinion in Biotechnology*, 19, 389–395.
- Doyle, J. C., & Csete, M. (2011). Architecture, constraints, and behavior. Proceedings of the National Academy of Sciences, 108(supplement_3), 15624–15630.
- Eckmann, J.-P., & Tlusty, T. (2021). Dimensional reduction in complex living systems: Where, why, and how. *BioEssays*, 43(9), 2100062.
- Frank, S. A. (2007). Maladaptation and the paradox of robustness in evolution. *PLoS ONE*, 2, e1021.

- Frank, S. A. (2017). Puzzles in modern biology. V. Why are genomes overwired? *F1000Research*, 6, 924.
- Frank, S. A. (2022a). Microbial life history: The fundamental forces of biological design. Princeton University Press.
- Frank, S. A. (2022b). Optimization of transcription factor genetic circuits. *Biology*, 11, 1294.
- Frank, S. A. (2023a). An enhanced transcription factor repressilator that buffers stochasticity and entrains to an erratic external circadian signal. *Frontiers in Systems Biology*, *3*, 1276734.
- Frank, S. A. (2023b). Precise traits from sloppy components: perception and the origin of phenotypic response. *Entropy*, 25, 1162.
- Frank, S. A. (2023c). Robustness and complexity. *Cell Systems*, 14(12), 1015–1020.
- Frank, S. A. (2023d). Robustness increases heritability: implications for familial disease. *Evolution*, 77, 655–659. qpac026.
- Frank, S. A. (2024a). A biological circuit to anticipate trend. *Evolution Letters*, 8, qrae027.
- Frank, S. A. (2024b). Circuit design in biology and machine learning. I. Julia software code.
- Friedlander, T., Mayo, A. E., Tlusty, T., & Alon, U. (2015). Evolution of bow-tie architectures in biology. *PLoS Computational Biology*, 11(3), e1004055.
- Gamoran, A., Lieberman, L., Gilead, M., Dobbins, I. G., & Sadeh, T. (2024). Detecting recollection: Human evaluators can successfully assess the veracity of others' memories. *Proceedings of the National Academy of Sciences*, 121(22), e2310979121.
- Gauthier, D. J., Bollt, E., Griffith, A., & Barbosa, W. A. (2021). Next generation reservoir computing. *Nature Communications*, 12(1), 1–8.
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press.
- Goudarzi, A., Lakin, M. R., & Stefanovic, D. (2013). Dna reservoir computing: A novel molecular computing approach. In Soloveichik, D., & Yurke, B. (Eds.), DNA computing and molecular programming (pp. 76–89). Springer International Publishing.
- Gray, M. W., Lukeš, J., Archibald, J. M., Keeling, P. J., & Doolittle, W. F. (2010). Irremediable complexity? *Science*, 330(6006), 920–921.
- Gupta, A., & Khammash, M. (2023). The internal model principle for biomolecular control theory. *IEEE Open Journal of Control Systems*, 2, 63–69.
- Gutenkunst, R. N., Waterfall, J. J., Casey, F. P., Brown, K. S., Myers, C. R., & Sethna, J. P. (2007). Universally sloppy parameter sensitivities in systems biology models. *PLoS Computational Biology*, 3, e189.
- Hebb, D. O. (1949). The organization of behavior. Wiley.
- Hiscock, T. W. (2019). Adapting machine-learning algorithms to design gene circuits. *BMC Bioinformatics*, 20, 214.
- Irie, N., & Kuratani, S. (2014). The developmental hourglass model: A predictor of the basic body plan? *Development*, 141, 4649–4655.
- Jaeger, H. (2007). Echo state network. Scholarpedia, 2(9), 2330.
- Jiahui, G., Feilong, M., Visconti di Oleggio Castello, M., Nastase, S. A., Haxby, J. V., & Gobbini, M. I. (2023). Modeling naturalistic face processing in humans with deep convolutional neural networks. *Proceedings of the National Academy of Sciences*, 120(43), e2304085120.
- Jordan, D. J., & Miska, E. A. (2023). Canalisation and plasticity on the developmental manifold of *Caenorhabditis elegans*. *Molecular Systems Biology*, 19(11), e11835.
- Kaneko, K. (2024). Constructing universal phenomenology for biological cellular systems: An idiosyncratic review on evolutionary dimensional reduction. *Journal of Statistical Mechanics: Theory* and Experiment, 2024(2), 024002.
- Kanwisher, N., Khosla, M., & Dobs, K. (2023). Using artificial neural networks to ask 'why' questions of minds and brains. *Trends in Neurosciences*, 46(3), 240–254.
- Kemp, C., & Tenenbaum, J. B. (2008). The discovery of structural form. Proceedings of the National Academy of Sciences, 105(31), 10687–10692.

- Kochanowski, K., Gerosa, L., Brunner, S. F., Christodoulou, D., Nikolaev, Y. V., & Sauer, U. (2017). Few regulatory metabolites coordinate expression of central metabolic genes in *Escherichia coli*. *Molecular Systems Biology*, 13(1), 903.
- Levin, M. (2024). Self-improvising memory: A perspective on memories as agential, dynamically reinterpreting cognitive glue. *Entropy*, 26(6), 481.
- Li, Q., Wu, Z., Ling, R., Feng, L., & Liu, K. (2020). Multi-reservoir echo state computing for solar irradiance prediction: A fast yet efficient deep learning approach. *Applied Soft Computing*, 95, 106481.
- Liu, H., Li, Z., Hall, D., Liang, P., & Ma, T. (2024). Sophia: A scalable stochastic second-order optimizer for language model pre-training.
- Loeffler, A., Zhu, R., Hochstetter, J., Diaz-Alvarez, A., Nakayama, T., Shine, J. M., & Kuncic, Z. (2021). Modularity and multitasking in neuro-memristive reservoir networks. *Neuromorphic Computing* and Engineering, 1(1), 014003.
- Loeffler, A., Zhu, R., Hochstetter, J., Li, M., Fu, K., Diaz-Alvarez, A., Nakayama, T., Shine, J. M., & Kuncic, Z. (2020). Topological properties of neuromorphic nanowire networks. *Frontiers in Neuroscience*, 14, 184.
- Lynch, M. (2007). The origins of genome architecture. Sinauer Associates.
- Lynch, M. (2012). Evolutionary layering and the limits to cellular perfection. Proceedings of National Academy Sciences USA, 109, 18851–18856.
- Lynch, M. R. (2024). Evolutionary cell biology: The origins of cellular architecture. Oxford University Press.
- Ma, F., & Zheng, C. (2023). Transcriptome age of individual cell types in Caenorhabditis elegans. Proceedings of the National Academy of Sciences, 120(9), e2216351120.
- Maass, W., Natschläger, T., & Markram, H. (2002). Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation*, 14(11), 2531–2560.
- Mao, J., Griniasty, I., Teoh, H. K., Ramesh, R., Yang, R., Transtrum, M. K., Sethna, J. P., & Chaudhari, P. (2024). The training process of many deep networks explores the same low-dimensional manifold. *Proceedings of the National Academy of Sciences*, 121(12), e2310002121.
- Martinuzzi, F., Rackauckas, C., Abdelrehim, A., Mahecha, M. D., & Mora, K. (2022). Reservoircomputing.jl: An efficient and modular library for reservoir computing models. arXiv, 2204.05117.
- McCulloch, W. S., & Pitts, W. (1943). A logical calculus of ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5, 115–133.
- Miikkulainen, R. (2025). Neuroevolution insights into biological neural computation. *Science*, 387(6735), eadp7478.
- Mishra, D., Ngam, P. N., & Tyo, K. E. (2023). Machine learning guided design and characterization of synthetic biology parts and circuits. *Metabolic Engineering*, 75, 86–101.
- Mitchell, M. (2023). AI's challenge of understanding the world. *Science*, 382(6671), eadm8175.
- Muñoz-Gómez, S. A., Bilolikar, G., Wideman, J. G., & Geiler-Samerotte, K. (2021). Constructive neutral evolution 20 years later. *Journal of Molecular Evolution*, 89, 172–182.
- Nielsen, A. A., Der, B. S., Shin, J., Vaidyanathan, P., Paralanov, V., Strychalski, E. A., Ross, D., Densmore, D., & Voigt, C. A. (2016). Genetic circuit design automation. *Science*, 352(6281), aac7341.
- Oda, K., & Kitano, H. (2006). A comprehensive map of the tolllike receptor signaling network. *Molecular Systems Biology*, 2(1), 2006–0015.
- O'Gieblyn, M. (2024). The trouble with reality. New York Review of Books, March 21, 2024.
- Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C., & Fried, I. (2005). Invariant visual representation by single neurons in the human brain. *Nature*, 435(7045), 1102–1107.

- Ramezanian-Panahi, M., Abrevaya, G., Gagnon-Audet, J.-C., Voleti, V., Rish, I., & Dumas, G. (2022). Generative models of brain dynamics. *Frontiers in Artificial Intelligence*, 5, 807406.
- Reddy, G. T., Reddy, M. P. K., Lakshmanna, K., Kaluri, R., Rajput, D. S., Srivastava, G., & Baker, T. (2020). Analysis of dimensionality reduction techniques on big data. *IEEE Access*, 8, 54776–54788.
- Rodan, A., & Tino, P. (2010). Minimum complexity echo state network. *IEEE Transactions on Neural Networks*, 22(1), 131–144.
- Rodan, A., & Tiňo, P. (2012). Simple deterministically constructed cycle reservoirs with regular jumps. *Neural Computation*, 24(7), 1822–1852.
- Ruach, R., Ratner, N., Emmons, S. W., & Zaslaver, A. (2023). The synaptic organization in the caenorhabditis elegans neural network suggests significant local compartmentalized computations. *Proceedings of the National Academy of Sciences*, 120(3), e2201699120.
- Rumelhart, D. E., & McClelland, J. L. (1986). Parallel distributed processing: Explorations in the microstructure of cognition, Volume 1: Foundations. MIT Press.
- Sabrin, K. M., Wei, Y., van den Heuvel, M. P., & Dovrolis, C. (2020). The hourglass organization of the *Caenorhabditis elegans* connectome. *PLoS Computational Biology*, 16(2), e1007526.
- Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., & Bengio, Y. (2021). Toward causal representation learning. *Proceedings of the IEEE*, 109(5), 612–634.
- Seoane, L. F. (2019). Evolutionary aspects of reservoir computing. *Philosophical Transactions of the Royal Society B*, 374(1774), 20180377.
- Shook, D. R., & Keller, R. (2008). Morphogenic machines evolve more rapidly than the signals that pattern them: lessons from amphibians. Journal of Experimental Zoology Part B: Molecular and Developmental Evolution, 310(1), 111–135.
- Solé, R., & Seoane, L. F. (2022). Evolution of brains and computers: The roads not taken. *Entropy*, 24(5), 665.
- Sorrells, T. R., & Johnson, A. D. (2015). Making sense of transcription networks. Cell, 161(4), 714–723.
- Sorzano, C. O. S., Vargas, J., & Montano, A. P. (2014). A survey of dimensionality reduction techniques.
- Stuber, G. D. (2023). Neurocircuits for motivation. *Science*, 382(6669), 394–398.
- Tang, D., Kang, R., Coyne, C. B., Zeh, H. J., & Lotze, M. T. (2012). Pamps and damps: Signals that spur autophagy and immunity. *Immunological Reviews*, 249(1), 158–175.
- Tishby, N., & Polani, D. (2011). Information theory of decisions and actions. In V. Cutsuridis, A. Hussain, & J. Taylor (Eds.), *Perception-action cycle* (pp. 601–636). Springer.
- Tripathi, S., Kessler, D. A., & Levine, H. (2023). Minimal frustration underlies the usefulness of incomplete regulatory network models in biology. *Proceedings of the National Academy of Sciences*, 120(1), e2216109120.
- Walleshauser, B., & Bollt, E. (2022). Predicting sea surface temperatures with coupled reservoir computers. Nonlinear Processes in Geophysics, 29(3), 255–264.
- Wu, C., Balakrishnan, R., Braniff, N., Mori, M., Manzanarez, G., Zhang, Z., & Hwa, T. (2022). Cellular perception of growth rate and the mechanistic origin of bacterial growth law. *Proceedings of the National Academy of Sciences*, 119(20), e2201585119.
- Wu, F., & Tan, C. (2023). Machine learning in synthetic biology: recent advances and future directions. *Nature Chemical Biology*, 19(6), 635–646.
- Yahiro, W., Aubert-Kato, N., & Hagiya, M. (2018). A reservoir computing approach for molecular computing. In ALIFE 2018: The 2018 Conference on Artificial Life (pp. 31–38). MIT Press.
- Zhao, J., Yu, H., Luo, J.-H., Cao, Z.-W., & Li, Y.-X. (2006). Hierarchical modularity of nested bow-ties in metabolic networks. BMC Bioinformatics, 7, 1–16.