



Learning to outgrow the competition: Reaction–diffusion systems that adapt to time-dependent environments

Olivier Rivoire^{a,1,2} and Guy Bunin^{b,1,2}

Edited by Anna C. Balazs, University of Pittsburgh, Pittsburgh, PA; received February 17, 2026; accepted July 16, 2026

A fundamental challenge in both biology and engineering is understanding how adaptation can emerge from simple physical or chemical building blocks. Biological adaptation appears to rely on two key capabilities: information processing, to enable learning of complex tasks, and reproduction, to give such learning its value through reproductive success. Creating and coordinating these capabilities poses both a formidable engineering challenge and a fundamental evolutionary puzzle. Here, we propose a model in which learning and reproduction emerge jointly from simple, nonequilibrium chemical systems, without prior design or evolution. The model's essential features are that it supports many steady states and spatial heterogeneities, upon which adaptation arises solely from exposure to time-varying influxes of reactants, without any predefined program or memory. The adaptive response is specific to the temporal sequence of environmental changes, yet general enough to extend to related sequences. Learning manifests in several forms: It is reinforced through repeated exposure to the same environmental sequence, representing a form of self-learning, and enhanced through spatial interactions, corresponding to a form of collective learning. Moreover, an adapted state can accelerate adaptation in nearby regions, providing a mechanism akin to teacher-guided learning. By coupling environmental variations to stochastic reaction–diffusion dynamics, our model establishes a minimal physical framework in which learning of complex environments and selective amplification of the states that encode them arise directly from nonequilibrium chemical processes.

adaptation | learning | reaction–diffusion

A hallmark of living systems is their ability to adapt to complex, changing environments. How does such behavior arise in biology, and how can it be engineered in artificial systems? While even the simplest organisms exhibit adaptation, they already represent highly intricate systems, with adaptive mechanisms that are far from fully understood (1). This drives efforts to find general principles of adaptation beyond contemporary biology, both to shed light on prebiotic origins of adaptive processes, and to guide the engineering of adaptive systems (2–6). Adopting a minimal, bottom–up perspective, the aim is to obtain rich, biological-like adaptation using only simple chemical or physical components, such as short molecules, small RNAs, or colloids (7–11). It is this significant challenge that we address in this work.

A fundamental difficulty is that biological adaptation seems to require two key abilities: information processing to enable adaptive responses to environmental stimuli and reproduction to express this adaptation through growth. Indeed, biological adaptation is ultimately defined by whether it leads to descendants that outgrow those of others. Considerable efforts in engineering and theory have gone into creating simple physical or chemical systems capable of learning (12–15), which involves information processing, and to designing simple self-replicators (10, 16–18). The two dominant frameworks for these goals are machine learning for learning and autocatalysis for replication. Yet each of these goals still presents challenges, and current frameworks offer no straightforward path to integrating the two capabilities. From a biological standpoint, the dominant framework to understand adaptation is evolution by natural selection, but it presents a chicken-and-egg conundrum: Since information processing and reproduction are prerequisites for evolution, they cannot have arisen through evolution itself.

Here, we propose a paradigm in which advanced learning and reproduction can emerge together from simple chemical systems, enabling complex adaptation. Unlike current frameworks, our approach does not rely on machine learning, autocatalysis, or natural selection, thereby overcoming their inherent limitations. Instead, our framework leverages systems with multiple nonequilibrium steady states and spatial heterogeneity, so that local regions can adopt distinct steady states and interact through diffusion (19).

Significance

Adaptive behavior in living systems relies on the interplay between learning and reproduction, yet it is unclear how these capabilities emerge and couple, starting from simple chemical processes. We introduce a minimal nonequilibrium chemical model in which learning and reproduction arise spontaneously, without prior design or Darwinian evolution. Relying on multistability and spatial heterogeneity, the system adapts solely through exposure to changing environmental conditions, in a specific yet robust manner. This manifests in multiple forms of learning, including reinforcement through repeated exposure and enhancement via spatial interactions. By demonstrating how selective amplification of better-performing states emerges from reaction–diffusion dynamics, this work provides a simple physical framework for the origin and design of adaptive behavior in chemical systems.

Author affiliations: ^aGulliver, CNRS, École supérieure de physique et de chimie industrielles de la ville de Paris, Université Paris Sciences et Lettres, Paris 75005, France; and ^bDepartment of Physics, Technion-Israel Institute of Technology, Haifa 32000, Israel

Author contributions: O.R. and G.B. designed research; performed research; and wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2026 the Author(s). Published by PNAS. This article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](#).

¹O.R. and G.B. contributed equally to this work.

²To whom correspondence may be addressed. Email: olivier.rivoire@espci.fr or buning@technion.ac.il.

Published August 24, 2026.

This framework does not involve explicit engineering of reproduction or learning, yet gives rise to states that can adapt in complex environments. Furthermore, it exhibits distinct modalities of adaptation, reminiscent of different forms of biological adaptation: self-adaptation through direct interaction with the environment, even in the absence of reproduction (analogous to physiological adaptation), and additional adaptation when reproduction through spatial expansion is enabled (analogous to evolutionary adaptation). This second modality can be interpreted as “collective learning” among a population of “students,” and is further enhanced by the presence of a “teacher.” Our paradigm thus offers both a fresh engineering perspective on chemical learning and fresh conceptual insights into biological adaptation.

A General Definition of Adaptation

We first present our definition of adaptation to time-varying environments in a general context before demonstrating how it arises in a particular system. This context is that of stochastic reaction–diffusion systems (Fig. 1A) described by an equation of the form

$$\partial_t n(x, t) = F(n(x, t), \epsilon(t)) + D \nabla^2 n(x, t) + \text{noise}(x, t). \quad [1]$$

Here $n(x, t)$ is an N -component field whose component $n_i(x, t)$ denotes the density of chemical species i at location x and time t ; $F(n(x, t), \epsilon(t))$ specifies the reaction kinetics, with a time-dependent parameter $\epsilon(t)$ representing the environment; $D \nabla^2 n(x, t)$ is the usual diffusion term; and $\text{noise}(x, t)$ accounts for stochasticity due to the finite number of molecules, as in chemical Langevin equations (see *Methods* for explicit expressions). We consider time-varying environments consisting of periodic drives $\epsilon(t)$ satisfying $\epsilon(t + \tau) = \epsilon(t)$ with a period τ .

To define adaptation, we propose a protocol consisting of two stages (Fig. 1B): a first stage in which the system is allowed to adapt and a second stage in which this adaptation is tested via competition. In the first (training) stage, the system experiences a particular periodic drive $\epsilon_1(t)$ for k_{train} periods, starting from an initial condition $n_0(x)$. We then select at random a location x^* and extract the final configuration at this location, $n_1 = n(x^*, k_{\text{train}}\tau)$. For comparison, the same system starting from the same initial condition is subjected to a different periodic drive $\epsilon_2(t)$ with the same period τ for the same total duration, and we retain the final state $n_2 = n(x^*, k_{\text{train}}\tau)$ at the same location x^* . In the second (testing) stage, we divide the spatial domain S into two equivalent subdomains S_1 and S_2 , and initialize each with one of the two previously obtained states, i.e. $n(x, 0) = n_1$ if $x \in S_1$ and $n(x, 0) = n_2$ if $x \in S_2$. We then run the dynamics for k_{test} periods under one of the two periodic drives, say $\epsilon_1(t)$. Finally, we say that the system has adapted to $\epsilon_1(t)$ if the final states $n(x, k_{\text{test}}\tau)$ for $x \in S_2$ are more similar to n_1 than to n_2 , on average over many realizations of the training and testing stages.

Model: Single Reactor

We will demonstrate that adaptation as defined in the previous section can arise in a specific system of reactions when subjected to a class of environmental sequences. First, we define this system and these environmental sequences.

Reaction Dynamics. We start by describing the reactions in the context of a single, well-mixed reactor. The key ingredients are N molecules A_i ($i = 1, \dots, N$) whose concentrations $[A_i]$ define the components n_i of the N -dimensional field n in Eq. 1. Each A_i




Fig. 1. (A) Our definition of adaptation applies to reaction–diffusion systems subjected to time-varying drives. Here, this is represented schematically, with reactions coupled by reciprocal inhibition in a series of reactors connected through diffusion. Time-varying influxes $\Lambda_i(t)$ of reactants define the environmental drive while uniform dilution (not represented) ensures that reaction volumes are kept constant. (B) Protocol for demonstrating adaptation. *Left:* First (training) stage, where the same system of reactions is subjected to two different periodic drives (here for over 10 cycles). The x-axis represents time, and the y-axis a one-dimensional space of coupled reactors. The densities $n(x, t)$ are colored based on their proximity to two of the final states, with red indicating closeness to a final state with the first environmental drive (*Top*), and blue to one with the second (*Bottom*). *Right:* Second (testing) stage, where these final states are competed against each other under repeated exposure to each environmental drive. Initially, half of the space is seeded with one final state and the other half with the other. After a few cycles (here 5), all reactors contain states close to the final state corresponding to their respective environmental sequence. When comparable results are obtained for most repetitions of the first and second stages, we say that the system adapts to the environmental drives.

is produced from a corresponding resource R_i when this resource is injected in the reactor. The N possible reactions of formation of A_i from R_i are coupled by sparse reciprocal inhibitions: When the product A_i of reaction i is present at high concentration, it prevents some other reactions j from occurring, and therefore the associated A_j from being produced. The simplest case with $N = 2$ reactions is illustrated in Fig. 2A. One possible mechanism is for the A_i to inhibit the catalyst required for forming A_j from R_j (*Methods*). The inhibitions are sparse, meaning that each A_i inhibits only a few reactions j , and they are reciprocal meaning that whenever A_i inhibits reaction j , A_j inhibits reaction i (denoted as $\vdash$).

These reactions take place in an open reactor in which each reactant R_i is supplied with an influx Λ_i and where constant dilution ensures a fixed volume. At any given time, each reactant R_i is either supplied ($\Lambda_i = 1$) or not supplied ($\Lambda_i = 0$). In the simple case with $N = 2$ reactions (Fig. 2A), when both R_1 and R_2 are supplied, corresponding to the influx vector $\Lambda = (1, 1)$, two stable steady states are possible: one where A_1 is high in concentration and A_2 is low and another where this is reversed.




Fig. 2. Systems, states, and environmental sequences. A system is specified by N nodes, with mutual inhibitions between some pairs of nodes, represented by --- . The environment Λ and the state of the system σ are N -dimensional vectors with elements in $\{0, 1\}$. Dynamically, a node becomes active, $\sigma_i = 1$, when its influx is turned on, $\Lambda_i = 1$, and no other node that inhibits it is active; a node becomes inactive as soon as its influx turns off, to $\Lambda_i = 0$. (A) A system with $N = 2$ nodes, with an inhibitory interaction. When all influxes are on, i.e., $\Lambda = (1, 1)$, the system can be in two states, $\sigma = (1, 0)$ or $\sigma = (0, 1)$. Top graph: continuous dynamics, showing stochastic switches between the two states; Bottom: the binary model used here. As indicated in the table where active nodes are represented in black and inactive ones in white, turning off one of the fluxes eliminates the possibility of one of the states. (B) The same principles apply to larger networks combining multiple nodes and interactions between them, where more states are present and where the relation between the influxes and the states is more intricate, as illustrated here with a network of $N = 4$ nodes. (C) An environmental sequence is defined by a series of $2m$ nodes $(i_1, \dots, i_{2m})$ where each node appears an even number of times. It represents a series of events where fluxes Λ_i are successively turned either on or off, starting and ending with all fluxes on. For example, in the environmental sequence $(1, 2, 1, 2)$, Λ_1 is turned off, followed by Λ_2 , before Λ_1 is turned on, followed by Λ_2 . Here, the dynamics of the system with $N = 2$ is represented when experiencing two particular environmental sequences, $(1, 2, 1, 2)$ and $(1, 2, 2, 1)$, starting each time from the two possible initial conditions, $\sigma = (1, 0)$ and $\sigma = (0, 1)$. In this example, the final state depends on the environmental sequence while the initial state is forgotten.

These concentrations are generally continuous but here we adopt a binary description valid at low noise (Fig. 2A and Methods) where one state is represented by $\sigma = (1, 0)$ and the other by $\sigma = (0, 1)$. If instead only R_1 is supplied, corresponding to $\Lambda = (1, 0)$, only the first state $\sigma = (1, 0)$ is possible (see table in Fig. 2A).

This generalizes to larger networks of reactions, represented by a graph whose nodes are the reactions i and whose edges indicate reciprocal inhibitions (Fig. 2B). When N is large, a very large number of steady states is possible. In the case of a single well-mixed reactor of large volume where stochastic fluctuations are negligible, these stable steady states are characterized by a vector of binary variables $\sigma \in \{0, 1\}^N$ that satisfy

$$(1 - \Lambda_i)\sigma_i = 0, \quad \forall i \quad \text{and} \quad \sigma_i\sigma_j = 0, \quad \forall i \text{---} j. \quad [2]$$

The first constraint imposes that $\sigma_i = 0$ whenever $\Lambda_i = 0$, and the second that σ_i and σ_j cannot both be 1 when reactions i and j reciprocally inhibit each other.

Noise due to the limited number of particles in reactors of finite volume can cause switches between the possible steady states. While the dynamics of these switches is fundamentally out-of-equilibrium, it is formally equivalent to the equilibrium dynamics of the thermal system with energy

$$E[\sigma] = - \sum_{i=1}^N \sigma_i \quad [3]$$

subject to the conditions of Eq. 2, and with temperature $T = 1/\sqrt{\Omega}$, where Ω is the finite volume of the reactor. This equivalence is permitted by the reciprocity of the inhibitions and is valid under a low-noise condition (Methods).

We introduce time-dependent drives by varying the influxes Λ_i . The time between two changes is chosen to be long enough to allow the system to relax to the nearest stable steady state following each environmental shift, yet short compared to the timescale required for full relaxation of the entire system. The stochastic dynamics is simulated using a kinetic Metropolis Monte Carlo algorithm.

Time-Dependent Environmental Sequences. We consider environmental sequences where $\Lambda(t)$ changes by turning on or off one resource R_i at a time (Fig. 2C). Initially, all resources are supplied, with $\Lambda_i = 1$ for all i . The first step consists in turning off the supply of R_{i_1} by setting Λ_{i_1} to 0. Recursively, the k th step consists in setting Λ_{i_k} to 0 if it was 1, and to 1 if it was 0. We perform $2m$ such steps, where each i_k appears an even number of times (possibly zero) so that after the last step, all the Λ_i have returned to 1. Thus, a given environmental sequence is represented by a series $(i_1, \dots, i_{2m})$ of reactions, where each reaction appears an even number of times. After each of the $2m$ steps, the system relaxes to the nearest stable state. These series of environmental changes generate transitions between states, as illustrated in Fig. 2C.

In the following figures, we consider $N = 50$ reactions with a pattern of reciprocal inhibition given by a random regular graph of connectivity 3. Unless otherwise mentioned, we show results averaged over 10 random graphs and, for each graph, over 1,000 independent trajectories, all starting with an empty state where $\sigma_i = 0$ for all i . Below, we generalize the model by considering multiple coupled reactors.

Encoding of Temporal Sequences

First, we look at states obtained after applying an environmental sequence, and study how these states encode information on the sequence. To focus on the encoding effects, we set $T = 0$ (no noise), so that when some Λ_i are turned on or off by setting them to 1 or 0, the state evolves to one of the nearest local minima of the energy function given by Eqs. 2 and 3. As only one Λ_i is turned on or off at a time, the dynamics are deterministic provided the graph of inhibitory interactions does not contain triangles, as we consider in this section.

We fix an environmental sequence with a given m and, starting from different initial conditions, record the states that we obtain after completing the environmental sequence. When $m/N \gg 1$, we obtain a single state independent of initial conditions (Fig. 3A). In this case, a large number of different final states is reached when sampling uniformly different environmental sequence. This number is exponential in N and not far from the total number of states when all $\Lambda_i = 1$ (Fig. 3B). In addition, the encoding is continuous, with similar environmental




Fig. 3. Encoding of information from environmental sequences in final states. (A) Mean distance between states reached after experiencing the same environmental sequence for varying values of m . Two trajectories are run under the same environmental sequence and the distance (normalized Hamming distance) between the two states at the end is reported. For m sufficiently large ($m/N \gg 1$) states obtained after the same environmental sequence are identical. *Inset:* mean number of active reactions (productivity) as a function of the number of steps with $m = 500$, starting from 3 different initial conditions, showing convergence. (B) Number of different states reached after a long environmental sequence ($m = 1,000$) when sampling uniformly environmental sequences, measured as the Shannon diversity, with an exponential fit, shown as a function of the number N of reactions. (C) Distance between states obtained from two environmental sequences differing by a given number of swaps, with $m = 1,000$. (D) Mean number of different state-cycles for a given environmental sequence (from exhaustive search).

sequences mapping, on average, to similar states: The distance between final states varies gradually with the number of swaps by which two environmental sequences differ, where one swap from environmental sequence $(i_1, \dots, i_{2m})$ consists in randomly choosing $1 \leq k < k' \leq 2m$ and swapping i_k and $i_{k'}$ (Fig. 3C). When $m/N \lesssim 1$, repeating successively the same environmental sequence causes the system to converge to periodic dynamics, repeating the same sequence of states with each environmental period. The number of such state-cycles associated with a given environmental sequence scales as $e^{N\phi(m/N)}$ where $\phi(m/N)$ decreases to 0 as m/N increases (Fig. 3D).

Competitive Advantage in Coupled Reactors

To analyze whether the information from past environments encoded in a state confers an adaptive advantage, we must quantify whether a state can reproduce more effectively than another. We enable reproduction by introducing spatial extension in the form of multiple reactors coupled by diffusion. This setup allows reproduction in the sense that a state initially confined to a subset of reactors can spread to other reactors. Reproduction here does not require strict replication; rather, it necessitates heredity whereby newly formed states must be correlated with their parent. Since a state is defined by which reactions are active or inhibited, this correlation is measured by the fraction of reactions that are similarly active or inhibited in both parent and offspring states. Using this capacity to reproduce, two states can be competed against each other by seeding them in neighboring reactors. If, after some time, the reactors each contain states more similar to the first than the second, the first state is considered to have reproduced at the expense of the second, and thus to be better adapted to the environment in which the competition occurred.

Model: Reaction-Diffusion Dynamics. Specifically, we consider coupling through diffusion $M > 1$ reactors arranged along a closed one-dimensional chain. The overall state of the system is represented by an $M \times N$ matrix $\sigma \in \{0, 1\}^{M \times N}$, whose dynamics are governed by an energy that generalizes Eq. 3,

$$E[\sigma] = - \sum_{x=0}^{M-1} \sum_{i=1}^N [\sigma_{x,i} + J\sigma_{x,i}\sigma_{(x+1),i}], \quad [4]$$

where $x+1$ is taken modulo M . The activation and inhibition constraints apply as before within each reactor x ,

$$(1 - \Lambda_i)\sigma_{x,i} = 0, \quad \forall x, i \quad \text{and} \quad \sigma_{x,i}\sigma_{x,j} = 0, \quad \forall x, i \vdash j. \quad [5]$$

The added spatial coupling term in Eq. 4 corresponds to diffusive dynamics in space; indeed, for continuous variables σ_i at temperature $T = 0$ where the dynamics are gradient descent in energy, this added energy term contributes to the dynamics a term of the form $d\sigma_{x,i}/dt = -\partial E/\partial \sigma_{x,i} = J[\sigma_{(x+1),i} + \sigma_{(x-1),i} - 2\sigma_{x,i}]$, which describes diffusion in discrete space (see *Methods* for discussion of the model with binary variables).

We start by considering two reactors coupled without periodic boundary conditions,

$$E[\sigma] = - \sum_{i=1}^N [\sigma_{1,i} + \sigma_{2,i} + J\sigma_{1,i}\sigma_{2,i}] \quad [6]$$

For large values of m , the environment uniquely determines the final states (Fig. 3A), making the outcome of competitions independent of previous experience. To avoid this trivial case of adaptation, we set $m = 40$, which allows for multiple state-cycles ($m/N = 0.8$ given $N = 50$, see Fig. 3D). We also set the duration of each environmental step to 10 units of time and the temperature to $T = 0.15$ to obtain approximately one state-cycle change every 10 environmental cycles that repeat a particular environmental sequence. Finally, we set $J = 0.5$.

Competition and Quantifying Adaptation. To assess the level of adaptation, we rely on competitions between states obtained after

Adaptive success




Fig. 4. Competition and adaptation. (A) Two states obtained in independent reactors under cycles of two distinct environmental sequences ϵ_1 and ϵ_2 are competed in two coupled reactors under cycles of environmental sequence ϵ_1 . Signed success is the fraction of times the state σ_3 resulting from the competition is more similar to σ_1 , the state evolved under the same environment, minus the fraction of cases where it is more similar to σ_2 . (B) The two environmental sequences ϵ_1 and ϵ_2 share a number of last steps in common. Given that there are $2m$ steps, they are identical when steps $m = 2$. (C) The environmental sequences of the two environments differ by n swaps, showing ability to discriminate changes in the order of steps within an environmental sequence.

Adaptive improvement




Fig. 5. Adaptive improvement. (A) Two competing states are obtained from the same trajectory at different times, a young state σ_1 and an old state σ_2 . (B) Here the old state is taken after 64 cycles, the young one from a varying number of cycles (horizontal axis).

cycles of different environmental sequences, as discussed before (Fig. 1). Given two environmental sequences ϵ_1 and ϵ_2 , we run c cycles of environmental sequence ϵ_1 in one isolated reactor, resulting in state σ_1 , and c cycles of environmental sequence ϵ_2 in another isolated reactor, resulting in state σ_2 , each an N -dimensional vector in $\{0, 1\}^N$. By concatenating these vectors into a $2 \times N$ matrix (σ_1, σ_2) , we construct the initial condition for a system with two coupled reactors, on which we run c' cycles of another environmental sequence ϵ'_1 , either identical to ϵ_1 , or related to it (as in Fig. 6B). This yields the pair of states (σ_4, σ_3) (Fig. 4A). By default, we set $c = 10^3$ and $c' = 10$.

Finally, we compare the similarity of σ_3 to σ_1 and to σ_2 , in order to assess the impact of an adapted state in a neighboring reactor. We define the signed difference as 1 if σ_3 is closer to σ_1 than to σ_2 in terms of the Hamming distance; -1 if it is closer to σ_2 ; and 0 if it is equally similar to both. This quantity is then averaged over reaction networks, environmental sequences, and stochastic trajectories to define a mean signed success.

Adaptation Without Reproduction

Assessing whether a state is adapted requires competition through reproduction, which in our model involves multiple reactors that are coupled by diffusion. Yet obtaining an adapted state does not necessarily require reproduction. As we now demonstrate, it can arise in a single well-mixed reactor, before its adaptation becomes manifest in coupled reactors.

Adaptive Success. We first consider competitions under cycles of an environmental sequence identical to that in which one of the competitors evolved, namely $\epsilon'_1 = \epsilon_1$. When ϵ_2 is unrelated to ϵ_1 , we observe that the signed success, averaged over different reaction networks, environmental sequences, and trajectories, is close to 1 (Fig. 4B, leftmost point). This indicates that the final state in reactor 2 is almost always more similar to the initial state in reactor 1 than to that in reactor 2. This observation also holds when the two environments ϵ_1 and ϵ_2 are more subtly different, such as when they share several last steps (Fig. 4B), or differ by a finite number of swaps (Fig. 4C).

Adaptive Improvement. One may wonder if the results shown in Fig. 4 A–C can be attributed to the response to the last

cycle only, independent of earlier environmental history. From a biological standpoint, one would like to differentiate between such a scenario, which is a basic form of plasticity, and learning, which is distinguished from other forms of plasticity by the fact that it improves upon repeated exposures. Here we demonstrate that such learning indeed happens by competing two states, σ_1 and σ_2 , taken along the same trajectory but differing in age (Fig. 5A). We find that, on average, the older state has an adaptive advantage over the younger one (Fig. 5B).

Generalization. From another perspective, that of machine learning, learning is a form of generalization. In this context, training involves examples drawn from a probability distribution, and testing other examples from the same distribution. To demonstrate this type of generalization in our model, we repeat the previous competition scheme (Fig. 6A), but now define a master environmental sequence ϵ_0 from which ϵ_1 and ϵ_3 are independently generated by performing a finite number of swaps, while ϵ_2 is obtained from ϵ_0 by an infinite number of swaps, representing an arbitrary permutation (Fig. 6B). The results show that adaptation does occur in this setting that requires generalization (Fig. 6C).

Enhanced Adaptation Through Interactions

Coupling to an Adapted State. To further demonstrate that adaptation occurs beyond environmental response, we analyze the effect of interacting with a reactor containing a previously adapted state (Fig. 7A). This teacher-student setting is implemented by a scheme where a “teacher” is first subjected to a number of environmental cycles, resulting in a state σ_0 . This state seeds one of two coupled reactors, while the other is initially empty, corresponding to the initial condition $(\sigma_0, 0)$. The final state in this initially empty reactor defines a student σ_1 that benefited from the teacher. As a control, a second student σ_2 is defined that does not benefit from a teacher, corresponding to the initial condition $(0, 0)$. Finally, we seed two coupled reactors with (σ_1, σ_2) to evaluate whether interacting with a teacher confers an adaptive advantage. The results confirm that this is indeed the case, provided teaching occurs over a sufficiently long but not too long time (Fig. 7B). The advantage of teaching eventually

Generalization




Fig. 6. Generalization. (A) Same scheme as Fig. 4A but (B) the environmental sequences applied during the competition differs from the two used in the training phase. The environmental sequences ϵ_1 and ϵ_3 are independently derived by n swaps from a reference ϵ_0 , while the environmental sequence ϵ_2 is an independent random permutation of ϵ_0 . (C) Generalization success as a function of n/m .




Fig. 7. Coupling to an adapted state. (A) A teacher is defined as a state that went through cycles of an environmental sequence. This state seeds a reactor that is coupled to another reactor. After a varying number of cycles in the same environmental sequence [x-axis in (B)], the final state of the second reactor defines an instructed student. Similarly, an untaught student is defined in the absence of a teacher. Finally, the instructed and uninstructed students compete in two coupled reactors. (B) Advantage in competition of being instructed as a function of the number of cycles in the second phase of the protocol. Interacting with a teacher is only beneficial if the training time is short enough so that a student without instruction does not entirely adapt on its own.

decreases over time, because with enough time even a state that has not been instructed can adapt perfectly.

Coupling Between Unadapted Systems. We now generalize to $M > 2$ coupled reactors arranged in a closed chain, as defined in Eqs. 4 and 5. We first assess whether interacting reactors provide an advantage by pitting two states, σ_1 and σ_2 , originating from systems with different numbers of coupled reactors, M and $M' \leq M$ respectively (Fig. 8A). Competition is set in a system of M reactors by seeding $M/2$ consecutive reactors with σ_1 and the remainder with σ_2 . Results demonstrate that larger systems yield better adapted states (Fig. 8B).




Fig. 8. Coupling between unadapted systems. (A) Competition in a system with M reactors between a state σ_1 obtained in a system of M reactors and a state σ_2 obtained in a system of $M' \leq M$ reactors. (B) Results reporting that, on average, the state obtained in the context of more reactors wins ($M = 32$ and varying $M' \leq M$). (C) Competition in two coupled reactors between a state obtained in the context of M reactors and a state obtained in a single reactor. (D) Results reporting that, on average, the state obtained with more reactors wins (varying M).

To further illustrate that (natural) selection under competition confers advantages beyond adaptation in a single reactor, we consider competition between a state obtained in multiple coupled reactors and one obtained in a single reactor, now with competition set in a system of $M = 2$ reactors (Fig. 8C). The former generally outperforms the latter, further confirming that interactions between reactors promote adaptation (Fig. 8D).

Conclusion

Designing artificial systems that emulate the remarkable adaptive capabilities of living systems in simplified physical or chemical setups is a long-standing goal in science and engineering. The elaborate information processing seemingly required for complex adaptation has led researchers to draw inspiration from computers and modern artificial intelligence. This influence plays an important role in the field of physical learning, where neural network architectures are commonly taken as a starting point that is then simplified to permit physical implementation (12). However, alternative principles may prove more relevant for systems far simpler than the brain.

Here, we propose such an alternative perspective, inspired by biological adaptation in organisms without nervous systems. Even microorganisms lacking identifiable information-processing organs are indeed capable of adaptive behavior: After exposure to a new environment, they can outperform their peers not exposed to this environment when competing under the same conditions (1, 20, 21). This observation leads to the speculation that the physicochemical processes underlying reproduction might already be sufficient, by themselves, to enable an elementary form of learning. Our model not only demonstrates this possibility but shows that they may already give rise to remarkably elaborate forms of learning.

Our approach builds on a minimal model of coupled chemical reactions that can be realized using only elementary rigid blocks that associate and dissociate under thermal fluctuations (19). These reactions generate a large variety of out of equilibrium states which differ by the combination of molecular concentrations they maintain. This is achieved through a network of reactions with sparse, random, and reciprocal inhibitions. As we showed previously, such dynamics reproduce core biological features: metabolism, heredity, and reproduction (19). Metabolism corresponds to maintaining specific out-of-equilibrium states driven by environmental inputs. Different influxes of reactants lead to distinct stable states, analogous to how a cell adopts different metabolic phenotypes depending on available nutrients. Heredity arises because stochastic perturbations, whether due to environmental fluctuations or molecular noise, can trigger transitions between states. Successive states being correlated, this creates a form of memory. Finally, reproduction refers to the ability of certain states to occupy increasingly larger parts of space over time, as they spread and expand at the expense of others.

Importantly, nothing in this setting is explicitly designed for learning, in particular by including the standard components of machine learning: There is no prescribed objective function, no a priori separation between “neurons” and “weights,” and no external supervision. Yet, when exposed to complex environmental sequences, the system spontaneously develops adaptive states, whereby the new states harbor a reproductive advantage under the conditions in which they emerged, mirroring adaptation in simple living organisms.

How does this adaptation arise? A well-known mechanism of biological adaptation is natural selection. Our system indeed

contains all ingredients necessary for Darwinian evolution by natural selection (19). However, we find that natural selection is not required for adaptation to occur as it can emerge in the absence of competition, in single well-mixed reactors where spatial heterogeneity, and thus reproduction, cannot take place. In this case, competition is only necessary to reveal adaptation: Competition between states shows that the system has reached states with enhanced reproductive potential tailored to specific environments. This intrinsic form of learning resembles physiological adaptation, which occurs within organisms. When spatial heterogeneity is introduced, the system also exhibits collective learning. This form of learning resembles evolutionary adaptation, which occurs between individuals of a population. Finally, our system also demonstrates educated learning through interaction with another system that previously experienced the environment.

The closest physical analogs to the present setting are disordered materials such as glasses, gels, granular materials, or mechanical metamaterials, whose hysteretic elements can encode and retain memory of past inputs (12, 22, 23). In well-mixed conditions, our model maps formally to a spin glass on a random graph used to describe some of these systems (*Methods*). Such materials, however, lack reproduction and thus do not display adaptation in the sense of reproductive advantage. Conversely, existing proposals for replicating soft-matter systems usually lack the internal diversity required to produce complex, stimulus-specific responses (24–26). Notable exceptions include DNA-based assemblies and multitype crystal systems that replicate and adapt through crystal growth and divisions (17, 27, 28). Yet, these designs rely on sophisticated engineering to reliably achieve replication. In contrast, our model exhibits adaptation without designing for it.

More broadly, while previous studies suggest that designing either learning or replication is a difficult task by itself and that coupling them may be even more challenging, our model shows that by focusing on spatially extended systems capable of supporting a large diversity of out-of-equilibrium states, reproduction, and learning can naturally arise and are intrinsically coupled. This result should help shed light on the remarkable adaptive abilities of simple organisms and suggest approaches for designing artificial systems with lifelike properties.

Methods

Here we describe how the present model with discrete variables and discrete space is related to the model with continuous variables and space described previously (19).

Continuous Model. The model consists of four types of molecular species, A_i , B_i , C_i , and R_i , with $i = 1, \dots, N$. The A_i are the primary species of interest whose relative concentrations define the system's state: They are produced from the resources R_i by means of the catalysts C_i ; can inhibit the catalysts C_j that mediate the formation of other A_j ; and this inhibition can be prevented when A_i binds to B_j . These reactions occur across three distinct timescales:



The second reaction only occurs between certain pairs, $i \mapsto j$, assumed to be reciprocally inhibitory, which are determined by the chemistry. Here we take, by design, that these pairs are sparse, with each i participating in three such pairs chosen at random.

The reactions take place in an open reactor driven by resource influx, where each R_i is injected at a given (possibly time-dependent) rate, and all species and complexes are diluted at a common rate δ , both defined on the slow timescale. The rates are chosen such that species C_i and B_i are effectively chemostated. Although summarized here at the level of chemical reactions, we showed before how these reactions can arise from simple building blocks that undergo Brownian motion and thermal association and dissociation (19).

Coarse-Graining. We first consider a single well-mixed reactor of large volume. The separation of timescales allows for coarse-graining into effective reactions that simplify the mathematical description. Starting on the fastest timescale with Eq. 9 and assuming strong interaction between A_i and B_j , the concentration of free A_i is the excess not bound to B_j , i.e. $[A_i]_{\text{free}} = [A_i]_{\text{tot}} - [B_j]_{\text{tot}}$ if $[A_i]_{\text{tot}} > [B_j]_{\text{tot}}$, and $[A_i]_{\text{free}} = 0$ otherwise, also written

$$[A_i]_{\text{free}} = \max(0, [A_i]_{\text{tot}} - [B_j]_{\text{tot}}) \quad [10]$$

Next, assuming that Eq. 8 equilibrates on an intermediate timescale, $[C_i A_j] = \Gamma_{ij} [C_i]_{\text{free}} [A_j]_{\text{free}}$, where Γ_{ij} is the association constant. As a consequence, the concentration of free C_i , which satisfies the conservation law

$$[C_i]_{\text{tot}} = [C_i]_{\text{free}} + \sum_{j \mapsto i} [C_i A_j] \quad [11]$$

is given by

$$[C_i]_{\text{free}} = \frac{[C_i]_{\text{tot}}}{1 + \sum_{j \mapsto i} \Gamma_{ij} [A_j]_{\text{free}}} \quad [12]$$

Finally, on the slowest timescale, combining Eq. 7 with the injection of R_i and the dilution of all complexes involving A_j , the total concentration of A_i follows

$$\partial_t [A_i]_{\text{tot}} = \alpha_i [C_i]_{\text{free}} [R_i] - \delta [A_i]_{\text{tot}} \quad [13]$$

All together, assuming all types to be equivalent with $\Lambda_i = \alpha_i [R_i] [C_i]_{\text{tot}}$, $\Gamma_{ij} = \Gamma$ and $[B_j]_{\text{tot}} = B$, the dynamics of the $n_i = [A_i]_{\text{tot}}$ is of the form

$$\partial_t n_i(t) = F_i(n(t), \Lambda) \quad [14]$$

with

$$F_i(n, \Lambda) = \Lambda_i c_i(n) - \delta n_i \quad [15]$$

and

$$c_i(n) = \frac{1}{1 + \Gamma \sum_{j \mapsto i} \max(0, n_j - B)}. \quad [16]$$

Out-of-Equilibrium Steady States. A simple example chemistry is $N = 2$ with reciprocal inhibition, $1 \mapsto 2$. This system admits two steady states, one obtained starting from a high concentration of A_1 which inhibits C_2 and keeps A_2 at low concentration; the other has the roles of A_1 and A_2 reversed. The role of B_i is to stabilize the inactive state of A_j .

For general N , up to 2^N combinations are possible, where each A_i is either inhibited or actively produced. However, not all are realizable steady states: If A_i and A_j inhibit each other, at most one can be active at a given time. Thus, steady states are characterized by a set of active nodes such that no two active nodes are connected, and no inactive node is connected exclusively to inactive nodes. These states correspond to solutions of a well-known graph-theoretic problem called independent sets, here applied to the inhibitory graph where each i is a node and where edges are defined by $i \mapsto j$. When each i is inhibited by three other j chosen at random, the number of steady states grows exponentially with N .

Diffusion. Space plays an important role in the model, as it enables growth. In a continuous setting, it is described by extending Eq. 14 to include diffusion,

$$\partial_t n_i(x, t) = F_i(n(x, t), \Lambda) + D \nabla^2 n_i(x, t) \quad [17]$$

Noise. Noise also plays an important role, as it enables random transitions between steady states. It is induced by fluctuations in the number of molecules. We describe it with chemical Langevin equations of the form of Eq. 1 with

$$\text{noise}(x, t) = T\sqrt{c_i(n(x, t)) + \delta n_i(x, t)} \xi_i(x, t) + T\sqrt{2D\partial_x} \left(\sqrt{n_i(x, t)} \zeta_i(x, t) \right), \quad [18]$$

where $\xi_i(x, t)$ and $\zeta_i(x, t)$ are Gaussian white noises, and where T is a “temperature” representing the strength of the noise.

Discrete Variables. At low enough noise, the dynamics in a single reactor switches between compositions resembling deterministic steady states (Fig. 2A). These states are described using discrete variables σ_i , where $\sigma_i = 1$ indicates that A_i is active and $\sigma_i = 0$ that it is inhibited.

We showed previously (19) that the out-of-equilibrium transitions between these discrete steady states satisfies detailed-balance, namely it is equivalent to equilibrium dynamics described by Metropolis rates, with an energy function

$$E[\sigma] = -\sum_{i=1}^N (\Lambda_i/\delta) \sigma_i \quad [19]$$

and the effective temperature $T = 1/\sqrt{\Omega}$, where Ω is the reactor volume. This energy is defined under the constraints that $\sigma_i = 0$ whenever $\Lambda_i = 0$, and $\sigma_i \sigma_j = 0$ whenever $i \mapsto j$ (the definition of realizable steady states). This mapping to an equilibrium process, which is not possible for arbitrary out-of-equilibrium dynamics, relies crucially on the reciprocity of inhibitory interactions.

Discrete Space. We also simplify the model by dividing space into distinct coupled reactors, where the concentrations in each reactor are uniform (well-mixed), and between the reactors elements diffuse at a rate ϕ . If this rate is small and if $\mathcal{S}(x)$ denotes the set of reactors coupled to reactor x , the reaction-diffusion equation, Eq. 17, becomes

$$\partial_t n_i(x, t) = \Lambda_i(t) c_i(n(x, t)) - \delta n_i(x, t) + \phi \sum_{y \in \mathcal{S}(x)} (n_i(y, t) - n_i(x, t)). \quad [20]$$

With discrete variables $\sigma_i(x)$ describing the state in each reactor x , and for small ϕ , transitions in each reactor effectively occur with a modified influx $\tilde{\Lambda}_i(x)$ that includes the diffusion from neighboring reactors $\mathcal{S}(x)$,

$$\tilde{\Lambda}_i(x) = \left(1 + \phi \sum_{y \in \mathcal{S}_x} (\sigma_i(y) - \sigma_i(x)) \right) \Lambda_i \quad [21]$$

When all reactors are coupled to the same number r of other reactors, the joint dynamics of states is governed by the energy

$$E[\sigma] = -(1 - r\phi)/\delta \sum_{i,x} \sigma_i(x) - \phi/\delta \sum_{i,x,y \in \mathcal{S}(x)} \sigma_i(x) \sigma_i(y) \quad [22]$$

This is the form that we use in Eq. 4, up to a multiplicative factor. The value used for J in our simulations, $J = 0.5$, is not so small, but results are qualitatively similar to the continuous model, on which Fig. 1B is based.

Spin-Glass Interpretation. The model in a single reactor, as defined by Eqs. 2 and 3, corresponds to a spin-glass model on a sparse graph, consisting of particles with hard-core repulsion. This model, known as the hard-core model, is also studied as a combinatorial optimization problem called vertex cover. On random graphs with low connectivity, it is known to exhibit a continuous (full-replica symmetry-breaking) glass transition at low temperature (29). From this point of view, Eq. 6 describes real replicas of hard-core models coupled along a one-dimensional line. To our knowledge, this spatial extension has not been previously studied in the spin-glass literature.

Data, Materials, and Software Availability. All study data are included in the main text.

- S. J. Gershman, P. E. Balbi, C. R. Gallistel, J. Gunawardena, Reconsidering the evidence for learning in single cells. *eLife* **10**, e61907 (2021).
- N. Goldenfeld, C. Woese, Life is physics: Evolution as a collective phenomenon far from equilibrium. *Annu. Rev. Condens. Matter Phys.* **2**, 375–399 (2011).
- C. Furusawa, K. Kaneko, Adaptation to optimal cell growth through self-organized criticality. *Phys. Rev. Lett.* **108**, 208103 (2012).
- N. Perunov, R. A. Marsland, J. L. England, Statistical physics of adaptation. *Phys. Rev. X* **6**, 021036 (2016).
- S. Landmann, C. M. Holmes, M. Tikhonov, A simple regulatory architecture allows learning the statistical structure of a changing environment. *eLife* **10**, e67455 (2021).
- K. Trifonova *et al.*, Trainable computation in molecular networks. *bioRxiv* [Preprint] (2025). <https://doi.org/10.64898/2025.12.28.696421> (Accessed 17 February 2026).
- J. M. Lehn, Perspectives in chemistry-aspects of adaptive chemistry and materials. *Angew. Chem. Int. Ed.* **54**, 3276–3289 (2015).
- H. W. van Roekel *et al.*, Programmable chemical reaction networks: Emulating regulatory functions in living cells using a bottom-up approach. *Chem. Soc. Rev.* **44**, 7465–7483 (2015).
- Z. Zeravcic, V. N. Manoharan, M. P. Brenner, Colloquium: Toward living matter with colloidal particles. *Rev. Mod. Phys.* **89**, 031001 (2017).
- P. Adamski *et al.*, From self-replication to replicator systems EN route to de novo life. *Nat. Rev. Chem.* **4**, 386–403 (2020).
- P. Pavlinova, C. N. Lambert, C. Malaterre, P. Nghe, Abiogenesis through gradual evolution of autocatalysis into template-based replication. *FEBS Lett.* **597**, 344–379 (2023).
- M. Stern, A. Murugan, Learning without neurons in physical systems. *Annu. Rev. Condens. Matter Phys.* **14**, 417–441 (2023).
- R. T. Nagipogu, D. Fu, J. H. Reif, A survey on molecular-scale learning systems with relevance to DNA computing. *Nanoscale* **15**, 7676–7694 (2023).
- J. D. Paulsen, N. C. Keim, Mechanical memories in solids, from disorder to design. *Annu. Rev. Condens. Matter Phys.* **16**, 61–81 (2025).
- K. M. Cherry, L. Qian, Supervised learning in DNA neural networks. *Nature* **645**, 639–647 (2025).
- M. Sipper, Fifty years of research on self-replication: An overview. *Artif. Life* **4**, 237–257 (1998).
- R. Schulman, B. Yurke, E. Winfree, Robust self-replication of combinatorial information via crystal growth and scission. *Proc. Natl. Acad. Sci. U.S.A.* **109**, 6405–6410 (2012).
- R. Zhuo *et al.*, Litters of self-replicating origami cross-tiles. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 1952–1957 (2019).
- G. Bunin, O. Rivoire, Evolutionary features in a minimal physical system: Diversity, selection, growth, inheritance, and adaptation. *Proc. Natl. Acad. Sci. U.S.A.* **122**, e2425753122 (2025).
- I. Tagkopoulos, Y. C. Liu, S. Tavazoie, Predictive behavior within microbial genetic networks. *Science* **320**, 1313–1317 (2008).
- A. Mitchell *et al.*, Adaptive prediction of environmental changes by microorganisms. *Nature* **460**, 220–224 (2009).
- N. C. Keim, J. D. Paulsen, Z. Zeravcic, S. Sastry, S. R. Nagel, Memory formation in matter. *Rev. Mod. Phys.* **91**, 035002 (2019).
- C. M. Meulblok, A. Singh, M. Labousse, M. van Hecke, Path-dependency and emergent computing under vectorial driving. *arXiv* [Preprint] (2025). <https://arxiv.org/abs/2503.07764> (Accessed 17 February 2026).
- L. S. Penrose, Self-reproducing machines. *Sci. Am.* **200**, 105–117 (1959).
- G. von Kiedrowski, A self-replicating hexadeoxynucleotide. *Angew. Chem. Int. Ed. Engl.* **25**, 932–935 (1986).
- Y. Sakref, O. Rivoire, Design principles, growth laws, and competition of minimal autocatalysts. *Commun. Chem.* **7**, 239 (2024).
- T. Wang *et al.*, Self-replication of information-bearing nanoscale patterns. *Nature* **478**, 225–228 (2011).
- R. Schulman, E. Winfree, How crystals that sense and respond to their environments could evolve. *Nat. Comput.* **7**, 219–237 (2008).
- J. Barbier, F. Krzakala, L. Zdeborová, P. Zhang, The hard-core model on random graphs revisited. *J. Phys. Conf. Ser.* **473**, 012021 (2013).