



Universiteit
Leiden
The Netherlands

Learning capacity and function of stochastic reaction networks

Ramezanpour, A.; Mashaghi, A.

Citation

Ramezanpour, A., & Mashaghi, A. (2023). Learning capacity and function of stochastic reaction networks. *Journal Of Physics: Complexity*, 4(3). doi:10.1088/2632-072X/acf264

Version: Publisher's Version

License: [Creative Commons CC BY 4.0 license](https://creativecommons.org/licenses/by/4.0/)

Downloaded from: <https://hdl.handle.net/1887/4082410>

Note: To cite this publication please use the final published version (if applicable).

PAPER • OPEN ACCESS

Learning capacity and function of stochastic reaction networks

To cite this article: Abolfazl Ramezanpour and Alireza Mashaghi 2023 *J. Phys. Complex.* **4** 035006

View the [article online](#) for updates and enhancements.

You may also like

- [Hyper-diffusion on multiplex networks](#)
Reza Ghorbanchian, Vito Latora and
Ginestra Bianconi
- [Higher-order connection Laplacians for
directed simplicial complexes](#)
Xue Gong, Desmond J Higham,
Konstantinos Zygalakis et al.
- [Sensitivity to network perturbations in the
randomized shortest paths framework:
theory and applications in ecological
connectivity](#)
Ilkka Kivimäki, Bram Van Moorter and
Marco Saerens



PAPER

OPEN ACCESS

RECEIVED
10 April 2023REVISED
8 July 2023ACCEPTED FOR PUBLICATION
21 August 2023PUBLISHED
1 September 2023

Original Content from
this work may be used
under the terms of the
[Creative Commons
Attribution 4.0 licence](#).

Any further distribution
of this work must
maintain attribution to
the author(s) and the title
of the work, journal
citation and DOI.



Learning capacity and function of stochastic reaction networks

Abolfazl Ramezanpour^{1,2} and Alireza Mashaghi^{2,3,*} ¹ Department of Physics, School of Science, Shiraz University, Shiraz 71946-84795, Iran² Medical Systems Biophysics and Bioengineering, Leiden Academic Centre for Drug Research, Faculty of Science, Leiden University, Leiden, The Netherlands³ Harvard Medical School, Boston, MA 02115, United States of America

* Author to whom any correspondence should be addressed.

E-mail: a.mashaghi.tabari@lacdr.leidenuniv.nl**Keywords:** stochastic reaction, learning, network topology

Abstract

Biochemical reaction networks are expected to encode an efficient representation of the function of cells in a variable environment. It is thus important to see how these networks do learn and implement such representations. The first step in this direction is to characterize the function and learning capabilities of basic artificial reaction networks. In this study, we consider multilayer networks of reversible reactions that connect two layers of signal and response species through an intermediate layer of hidden species. We introduce a stochastic learning algorithm that updates the reaction rates based on the correlation values between reaction products and responses. Our findings indicate that the function of networks with random reaction rates, as well as their learning capacity for random signal-response activities, are critically determined by the number of reactants and reaction products. Moreover, the stored patterns exhibit different levels of robustness and qualities as the reaction rates deviate from their optimal values in a stochastic model of defect evolution. These findings can help suggest network modules that are better suited to specific functions, such as amplifiers or dampeners, or to the learning of biologically relevant signal-response activities.

1. Introduction

Chemical transformations play a central role in the field of chemistry and encompass a vast array of diverse chemical reactions [1]. In living organisms, biochemical reaction networks are responsible for a wide range of biological functions, such as metabolism and signal transduction [2]. Decoding the design principles underlying these complex biochemical reactions is required for life-inspired reaction engineering and understanding the emergence of life on Earth.

A fundamental feature of biochemical reaction networks is the co-existence of both complexity and robustness in these self-organizing, self-regulating systems [3]. Homeostasis is a recurring theme in biology with a long history going back to Claude Bernard who emphasized the importance of maintaining ‘le milieu intérieur’ [4]. In biological organism, regulated variables often robustly adapt to environmental perturbations [5]. Furthermore, robustness is often a desired feature of engineered reaction networks [6]. It is thus important to see how chemical networks can learn and implement representations of a certain function.

Hebb’s rule is a natural mechanism of learning in (artificial) neural networks, especially in models of associative memory [7]. The back-propagation algorithm provides an efficient strategy when it comes to training of deep (multilayer) feed-forward networks [8, 9]. The function and learning performance of these networks are well studied and many powerful mathematical tools and concepts have been developed, e.g. see [10–19] and references therein. Regarding the reaction networks, we know that stochastic reaction networks are Turing universal and can compute any computable function for any nonzero error probability [20–22]. There are many interesting works which use chemical species and reactions as computational resources to implement the basic gates and circuits that are needed for a universal computation [23–26]. An important problem here is reconstruction of a reaction network or learning of the reaction rates which reproduce a set

of empirical or expected fluxes [27–39]. It is usually easier here to work with deterministic differential equations (or the steady states) in the thermodynamic limit, ignoring intrinsic fluctuations and correlations in the number of species for a finite system.

Currently, our understanding of learning capacity of stochastic chemical reactions is limited. Previous research into the function and learning of reaction networks has been mostly focused on deterministic models in the thermodynamic limit where the number of species and volume approach infinity but the concentrations are finite. In this limit, the dynamics of the system is described by a set of coupled ordinary differential equations for the average concentrations, assuming for instance the mass-action or Michaelis–Menten kinetics. A mathematical treatment of the space of solutions and functions that can be exhibited by these equations is presented in [40]. The computational power of deterministic reaction systems is demonstrated in [20]. Similarities with neural networks was exploited in [37–39] to have a mapping of the dynamics in the two systems and to implement different neural circuits with a reaction network. An elementary cellular signalling pathway is trained in [27] to display a specific signal-response pattern. The learning of linearly separable and inseparable problems is simulated in [30, 34] using only chemical species and reactions.

The first step to study the learning capacity of stochastic chemical reactions is to characterize the function and learning capabilities of basic artificial reaction networks. In this article, we are interested in stochastic reaction networks in finite systems where fluctuations are not negligible. These networks can compute any computable function with a nonzero error probability [21, 22]. The function of random neural network and deterministic reaction networks is studied in [14–16, 41]. Here we do a similar study for random networks of stochastic reactions. Efficient simulation of stochastic reaction networks is in general a computationally hard task but there are good approximate algorithms for simulating these systems [42–45]. Learning capacity of multilayered neural networks has been studied in [10–12] using the methods of statistical physics. In this work, we shall investigate the learning performance of stochastic reaction networks which has not been studied before as far as we are aware of. The stability or robustness of deterministic reaction networks in a changing environment are discussed in [4, 46, 47]. In [48] the difference between the robustness in deterministic and stochastic networks is illustrated. In this work, we see how much the stored patterns in a stochastic reaction network are robust to mutations in the reaction rates in a simple model of disease evolution introduced in [49].

Here, we focus on simple networks of reversible reactions and provide a proof-of-concept analysis of learning process in stochastic chemical reaction networks. We estimate the learning capacity and its dependency on reaction properties and network topology. We also assess the quality and robustness of the stored patterns. In the following, we start with the model definitions in section 2 and present the results of numerical simulations in section 3. The results are reported in three subsections dealing with the networks' functions and capacities, learning of random signal-response patterns, and robustness of the stored patterns. The concluding remarks are given in section 4. Overall, our finding provides insights relevant to engineering reaction networks with desired learning capacities.

2. Model and setting


2.1. Preliminaries

We consider a well-stirred mixture of N interacting species connected with M reversible chemical reactions (i.e. $2M$ irreversible reactions). We use $0 \leq X_i(\tau) \leq X_{\max}$ to show the number of molecules for species $i = 1, \dots, N$ at time τ . No conservation law for the number or mass of the molecules is assumed here. The maximum number $X_{\max}$ scales with the system volume V . Generally, a reaction $a = 1, \dots, 2M$ is represented by the stoichiometric coefficients $\nu_a^-(i)$ and $\nu_a^+(i)$

$$\sum_{i=1}^N \nu_a^-(i) \hat{X}_i \rightarrow \sum_{j=1}^N \nu_a^+(j) \hat{X}_j. \quad (1)$$

When the reaction occurs, the reactant/product species in the left/right hand side of the reaction are consumed/generated and X_i decreases/increases by $\nu_a^+(i) - \nu_a^-(i)$. We use ν_a^- and ν_a^+ for the set of reactants and products in reaction a , respectively.

In the following, we assume that the coefficients $\nu_a^-(i), \nu_a^+(i) \in \{0, 1\}$. Moreover, we shall consider reactions that connect at most two distinct reactant species to two other distinct product species. A reaction of type Rmn involves m reactants and n products. In the reaction networks that are studied here, we have only one-to-one (R11), one-to-two (R12), two-to-one (R21), or two-to-two (R22) reactions. A reaction network is called an Rmn network if all the reactions are of type Rmn .



The dynamics is modeled by a continuous-time Markov chain where each reaction has a chance to occur in a small time interval $d\tau$ with a given probability [50]. The propensity function $\eta_a(\mathbf{X})$ times $d\tau$ gives the probability that reaction a occurs in the time interval $(\tau, \tau + d\tau)$ given the number of molecules $\mathbf{X}(\tau) = \{X_1(\tau), \dots, X_N(\tau)\}$. To be more precise, the propensity function is solely dependent on the number of molecules located on the left-hand side of the reaction $\nu_a^- = \{i : \nu_a^-(i) > 0\}$. In the mass-action kinetics with $\nu_a^-(i), \nu_a^+(i) \in \{0, 1\}$, we have:

$$\eta_a(\mathbf{X}) = \kappa_a \prod_{i \in \nu_a^-} x_i, \quad (2)$$

where $x_i(\tau) = X_i(\tau)/X_{\max}$ and the reaction rates $\kappa_a \propto X_{\max}$. The latter is to ensure that the average concentrations $\langle x_i \rangle$ are well defined in the thermodynamic limit $V, X_{\max} \rightarrow \infty$. In practice, and in numerical simulations, we work with finite and fixed values for V and $X_{\max}$. Specifically, in the following we shall set $V = X_{\max}$.

2.2. Network

The definition of a reaction network comprises three components: the set of species or reactants (nodes), the reactions connecting the species (edges), and the reaction rates (couplings). In what follows, we shall consider multilayered reaction networks starting with one layer of signals and ending with one layer of responses. Figure 1 depicts a Petri net representation of the reaction networks. Assuming that the signal in a biochemical reaction system is encoded in the concentration of the signal species, let $\mathcal{S}$ be the set of signal variables which is a subset of species variables $\mathcal{I} = \{i : i = 1, \dots, N\}$. The signal variables are externally driven to maintain a fixed activity level that is independent of the other variables. This scenario can occur in homeostatic regulation of cellular reaction networks or when a biochemical reaction system is being manipulated with drugs. In addition, we have the disjoint subset of response variables $\mathcal{R} \subset \mathcal{I}$. These variables could be on any side of a reaction. The number of signal and response variables is denoted by $N_S = |\mathcal{S}|$ and $N_R = |\mathcal{R}|$, respectively.

Note that a (reversible) reaction in these networks only connects species in two consecutive layers. The signal and response species do not appear in the hidden layers and each species corresponds to a single node in the network. N_h determines the quantity of variables in the hidden layers. The hidden species are to process (e.g. filter or amplify) and transfer the signals to the response species (e.g. from surface receptors to nucleus of cell). As in artificial neural networks, the hidden layers of species are expected to enhance expressivity of multilayered reaction networks. We assume that the number of signals is larger than that of responses to model redundancy of pathways to the responses. Considering these points, in the following we work with the

largest networks of stochastic reactions that are computationally tractable to have a reasonable statistics. We shall see how the structure and the nature of reactions determine the function of such reaction networks.

2.3. Dynamics

To simulate the time evolution of a stochastic reaction network, the Stochastic Simulation Algorithm (or Gillespie algorithm) can be used [42]. Suppose that a reaction occurred at time τ . Then, we need to know the time τ_a to the next reaction and its index a . Here we assume that the time interval $d\tau$ is small enough such that at most one reaction can occur in that interval. This defines a Poisson process where the probability of occurring a reaction is also independent of the other reactions. Thus, assuming the system is in state $\mathbf{X}(\tau)$, the joint probability distribution of the two random variables (τ_a, a) can be expressed as follows:

$$\rho(\tau_a, a) = \eta_a(\mathbf{X})e^{-\eta(\mathbf{X})\tau_a}, \quad \eta(\mathbf{X}) = \sum_a \eta_a(\mathbf{X}). \quad (3)$$

To generate samples of the aforementioned random variables, the Gillespie algorithm operates as follows:

- compute the η_a and η given the $\mathbf{X}(\tau)$
- extract τ_a and a from the probability distribution $\rho(\tau_a, a)$:
 - * draw two independent random numbers u_1, u_2 from the uniform distribution in $(0, 1)$
 - * then $\tau_a = \frac{1}{\eta} \log(1/u_1)$ and a is the smallest integer satisfying $\sum_{b=1}^a \eta_b > u_2 \eta$
- change $\tau \rightarrow \tau + \tau_a$ and $\mathbf{X} \rightarrow \mathbf{X} + (\nu_a^+ - \nu_a^-)$

Note that the numbers of molecules for externally driven (signal) species are fixed and do not follow the above dynamics even if they are involved in a reaction. Moreover, we have this constraint that $0 \leq X_i \leq X_{\max}$ for all species. We take $X_{\max} = 1000$ in all the following numerical simulations. This is on the order of the number of molecules present in a cell. When $X_{\max}$ (and hence X_i) takes on very small values, large finite size fluctuations occur and the uncertainty in the number of molecules dominates. On the other hand, for very large values of $X_{\max}$ and extensive X_i , the relative importance of fluctuations around the mean values would be negligible and one recovers the results of deterministic models (differential equations) for the average concentrations. One expects to observe different qualitative behaviors in these two limits. Interesting behaviors are usually exhibited in the region separating the two regimes, where real cells function.

Given an initial state for the number of molecules $\mathbf{X}(0)$ and the magnitude of κ_a , the above algorithm generates a sequence of states where in each step a single reaction occurs. An iteration of dynamics is defined by $2M$ steps. $\mathbf{X}(t)$ represents the state of the system, indicating the number of molecules for each species, following t iterations. In the following, the number of iterations is chosen large enough to ensure that the system is in the stationary state of the dynamics. Then samples of $\mathbf{X}(t)$ are extracted to obtain, e.g. the average values of concentrations and correlations between the concentrations of two species. Throughout the following sections, this algorithm will be repeatedly employed to obtain the above statistical information for a given realization of the reaction rates.

2.4. Signal-response patterns

Let us consider a reaction network which displays responses $\mathbf{r}^\mu$ in presence of signals $\mathbf{s}^\mu$. Here $\{\mathbf{s}^\mu, \mathbf{r}^\mu\}$ denote the expected activation values of the signal and response variables for a set of P patterns indexed by $\mu = 1, \dots, P$. The signals have the specified values $s_i^\mu = 0, 1$ in pattern μ ; that is for pattern μ the average concentration satisfies $c_i^\mu > \frac{1}{2} + \delta$ if $s_i^\mu = 1$, and $c_i^\mu < \frac{1}{2} - \delta$ for $s_i^\mu = 0$. The interval 2δ defines the size of the uncertainty gap between the two states. Similarly, for the response variables $c_j^\mu > \frac{1}{2} + \delta$ means that $r_j^\mu = 1$ and $c_j^\mu < \frac{1}{2} - \delta$ is equivalent to $r_j^\mu = 0$. In the following, we set $\delta = 1/6$. Note that any species, including signals and responses, could be active or inactive depending on its concentration: it is active for concentrations greater than $\frac{1}{2} + \delta$ and inactive for concentrations less than $\frac{1}{2} - \delta$. For concentrations between the two limits the species are considered neutral. The uncertainty gap 2δ should be larger than the standard deviations in the concentrations to have a well defined active state and inactive state.

Later in this study, to investigate the function space of a random reaction network, we compute the responses (output) for all the possible signal activities (input), given a realization of the random reaction rates. The number of such signal-response patterns is 2^{N_s} . In the learning part of the study, however, we look for a stochastic reaction network which stores P randomly and uniformly chosen signal-response patterns. In both cases, the reaction network structure remains constant, and only the reaction rates κ_a vary.

3. Results

3.1. Network capacity and function

To characterize the function of a reaction network we study the statistics of response concentrations $\{x_j^\mu : j \in \mathcal{R}\}$ for all the possible 2^{N_s} signal patterns [14, 15]. The average concentration $c_j^\mu = \langle x_j^\mu \rangle_t$ denotes the time average of concentration in the stationary state for signal pattern $\mathbf{s}^\mu$. The results are obtained for an ensemble of random and independent reaction rates κ_a uniformly distributed in $(0, \kappa_{\max})$. This ensemble just provides the maximum range of models that here are investigated and it does not mean that in reality (or at equilibrium) the reaction rates are distributed uniformly and independently. That is we take a random sample of the reaction rates κ and run the Gillespie algorithm for signal activity $\mathbf{s}^\mu$ to obtain the time averages $c_j^\mu(\kappa)$. This process is repeated for different realizations of κ to estimate the ensemble averages over the space of the reaction rates. In the following, we will consider the time-averaged concentrations and correlations of the response variables for a signal pattern μ ,

$$c_R^\mu = \frac{1}{N_R} \sum_{j \in \mathcal{R}} \langle c_j^\mu(\kappa) \rangle_\kappa, \quad (4)$$

$$c_{RR}^\mu = \frac{2}{N_R(N_R - 1)} \sum_{(ij) \in \mathcal{R}} \langle c_{ij}^\mu(\kappa) \rangle_\kappa. \quad (5)$$

The time-averaged correlations are

$$c_{ij}^\mu(\kappa) = \frac{\langle x_i x_j \rangle_t - \langle x_i \rangle_t \langle x_j \rangle_t}{\sqrt{(\langle x_i^2 \rangle_t - \langle x_i \rangle_t^2)(\langle x_j^2 \rangle_t - \langle x_j \rangle_t^2)}}. \quad (6)$$

Note that the averages with respect to time and random realizations of the reaction rates are denoted by $\langle \cdot \rangle_t$ and $\langle \cdot \rangle_\kappa$, respectively.

Given a signal pattern μ , we are also interested in the average polarization of the responses

$$p_R^\mu = \frac{1}{N_R} \sum_{j \in \mathcal{R}} \left\langle \left| c_j^\mu(\kappa) - \frac{1}{2} \right| \right\rangle_\kappa, \quad (7)$$

to see how much the responses' concentrations are deviated from the neutral value $1/2$. The overlap of responses for two random and independent set of reaction rates is denoted by

$$q_R^\mu = \frac{1}{N_R} \sum_{j \in \mathcal{R}} \left\langle \left(c_j^\mu(\kappa) - \frac{1}{2} \right) \left(c_j^{\mu'}(\kappa') - \frac{1}{2} \right) \right\rangle_\kappa. \quad (8)$$

This quantity is a measure of changes in the responses that are induced by varying the reaction rates. We also check the overlap of responses for two different signal patterns μ and μ' , to have a measure of variation in the responses by the input patterns,

$$q_R^{\mu\mu'} = \frac{1}{N_R} \sum_{j \in \mathcal{R}} \left\langle \left(c_j^\mu(\kappa) - \frac{1}{2} \right) \left(c_j^{\mu'}(\kappa') - \frac{1}{2} \right) \right\rangle_\kappa. \quad (9)$$

To characterize the signal patterns, we define the level of activities and overlap of two signal patterns,

$$m_S^\mu = \frac{1}{N_S} \sum_{i \in \mathcal{S}} (2s_i^\mu - 1), \quad (10)$$

$$q_S^{\mu\mu'} = \frac{1}{N_S} \sum_{i \in \mathcal{S}} (2s_i^\mu - 1)(2s_i^{\mu'} - 1). \quad (11)$$

These quantities provide the basic statistics of the input signals. Note that m_S^μ is related to the number of active signals in pattern μ and is distinct from the mean concentration of a signal variable represented as $c_i^\mu(\kappa)$.

Let us start with the function of fully connected one-layer reaction networks with iid reaction rates κ_a uniformly distributed in $(0, \kappa_{\max})$. A fully connected network of Rmn reactions is a network in which all possible reactions of type Rmn are present between two consecutive layers of the network. In all the following numerical simulations we set $\kappa_{\max} = 0.01X_{\max}$. That is proportional to $X_{\max}$ but small enough to allow us to

Table 1. Response of the one/two-layer reaction networks. The average concentrations c_R^μ and correlations c_{RR}^μ of response variables in a fully connected network of $N_S = 10$ signals and $N_R = 5$ responses. The number of hidden variables is $N_h = 5$ in the two-layer network. The results are averaged over independent realizations of the reaction rates for all the possible signal patterns. The number of iterations is $t = 2000$ for the R11, R12, and $t = 4000$ for the R21, R22 networks.

Reactions	Concentration (one/two-layer)	Correlation (one/two-layer)
R11	$(0.523 \pm 0.002)/(0.53 \pm 0.003)$	$(-0.009 \pm 0.0008)/(-0.019 \pm 0.0008)$
R12	$(0.714 \pm 0.001)/(0.83 \pm 0.004)$	$(0.004 \pm 0.0009)/(0.028 \pm 0.003)$
R21	$(0.253 \pm 0.0004)/(0.06 \pm 0.001)$	$(-0.022 \pm 0.0007)/(-0.007 \pm 0.0003)$
R22	$(0.498 \pm 0.003)/(0.50 \pm 0.027)$	$(-0.028 \pm 0.001)/(-0.002 \pm 0.007)$

simulate the stochastic dynamics of the species in a reasonable computational time. We sample a random realization of the reaction rates from the above distributions and run the Gillespie algorithm for t iterations to reach the stationary state of the dynamics. The interested quantities are then sampled and the process repeated for different realizations of the reaction rates to obtain the statistics of the stationary state. We checked that the results are not qualitatively sensitive to small changes in the value of $\kappa_{\max}$, and the number of signals, responses, and hidden variables. This is expected to be true also for different distributions of the reaction rates with nearly the same means and variances. However, we needed more iterations to reach the stationary state of the networks with the R21 and R22 reactions ($t = 4000$) than with the R11 and R12 reactions ($t = 2000$).

Table 1 reports the average activity and correlation of the response variables in such networks. We observe that the activity is close to $1/2$ for both the R11 and R22 networks, whereas the R12 and R21 networks have high and low average activities, respectively. The response variables are nearly uncorrelated, ignoring the very small correlations of order ± 0.01 . Similar behaviors are observed in table 1 for the two-layer networks where the effects of the R12 and R21 networks are enhanced. This means that on average the latter networks work respectively as an amplifier and dampener irrespective of the inputs and the reaction rates. Moreover, one can enhance these effects by increasing the number of layers in the network. In fact, when the number of reactants is larger than the number of products the probability of direct reaction is smaller than that of the reverse one according to the mass action kinematics. That is why we see that for R21 reactions the concentration of responses is smaller than $1/2$ whereas for R12 reactions it is the opposite.

Figures 2 and 3 display the statistics of the signal and response variables in the same networks. In average, the response polarization p_R^μ and overlaps $q_R^\mu, q_R^{\mu\mu'}$ are larger for the R12 and R21 networks than for the R11 and R22 networks. That is the responses are closer to 0 and 1 and more similar to each other for different patterns and reaction rates. Moreover, in the R11 and R22 networks the p_R^μ and q_R^μ are nearly symmetric around the signal activity $m_S^\mu = 0$, where half the signals are active and the other half inactive. On the other hand, the two quantities p_R^μ and q_R^μ are increasing with m_S^μ in the R12 network and are decreasing with m_S^μ in the R21 network. Figures 2 and 4 show the above statistics for the two-layer reaction networks which are qualitatively similar to the one-layer networks. In words, in both the one- and two-layer networks the R12 and R21 reactions favor the ‘on’ and ‘off’ states, respectively. On the other hand, the two states are balanced in networks with the R11 and R22 reactions and one can switch between these states by changing the number of active signals. De Palma *et al* [16] demonstrates that random neural networks tend to exhibit bias towards simple functions. Interestingly, we observe that random reaction networks are very similar to random neural networks in supporting simple functions.

The learning capacity of a reaction network can be quantified as follows. Consider the ensemble of random and independent patterns $\{\mathbf{s}^\mu, \mathbf{r}^\mu\}$ and define the success probability $\Pr(P)$ of storing P of such patterns. The network capacity is defined by the threshold P_c where the success probability approaches zero, $\Pr(P > P_c) \rightarrow 0$. On the other hand, one can look at the number of solutions in the reaction space $\kappa = \{\kappa_a : a = 1, \dots, 2M\}$,

$$\mathcal{N}_\kappa = \int D\kappa \prod_\mu I_\mu(\kappa), \quad (12)$$

where $\int D\kappa$ denotes an integral over the reaction space $0 < \kappa_a < \kappa_{\max}$. The indicator function $I_\mu(\kappa) = 1$ if pattern μ is stored, otherwise $I_\mu(\kappa) = 0$. We say that pattern μ is stored if the signal values $\mathbf{s}^\mu$ result in the associated responses $\mathbf{r}^\mu$ by following the system dynamics. The capacity is given by the maximum number of patterns where the average entropy $\log \mathcal{N}_\kappa$ vanishes [12].

Consider a one-layer network of the R11 and R21 reactions with a single response variable ($N_R = 1$). Given the signals s_i^μ , the average concentration of the response variable in the thermodynamic limit c_j^μ is governed by the following deterministic equation




Figure 2. The response of the one/two-layer R11, R12, R21 and R22 reaction networks. The panels display the distribution (KDE) of the polarization p_R^μ , and the response overlaps q_R^μ and $q_R^{\mu\mu'}$ for (a1)–(c1) one-layer and (a2)–(c2) two-layer networks. The κ_a are chosen randomly and uniformly in $(0, 10)$ in fully connected networks of size $(N_S = 10, N_R = 5)$. The number of hidden variables $N_h = 5$ in the two-layer network. The number of iterations is $t = 2000$ for the R11, R12, and $t = 4000$ for the R21, R22 networks.

$$\frac{d}{dt}c_j^\mu = \sum_{a \in A_\downarrow} \kappa_a \prod_{i \in \nu_a^-} c_i^\mu - \sum_{a \in A_\uparrow} \kappa_a c_j^\mu, \quad \forall \mu. \quad (13)$$

Here the c_i^μ are the average concentrations of the signal variables constrained by $|c_i^\mu - s_i^\mu| < \frac{1}{2} - \delta$. Moreover, $A_\downarrow$ refers to the set of reactions that are directed from the upper layer to the lower layer. Similarly, we define the set $A_\uparrow$ of reverse reactions. In the stationary state,

$$0 = \sum_{a \in A_\downarrow} \kappa_a \prod_{i \in \nu_a^-} c_i^\mu - \sum_{a \in A_\uparrow} \kappa_a c_j^\mu, \quad \forall \mu. \quad (14)$$

A solution in the reaction space should satisfy the following constraints:

$$0 < c_j^\mu = \sum_{a \in A_\downarrow} \tilde{\kappa}_a \prod_{i \in \nu_a^-} c_i^\mu < 1, \quad (15)$$

$$r^\mu = 1 \rightarrow c_j^\mu = \sum_{a \in A_\downarrow} \tilde{\kappa}_a \prod_{i \in \nu_a^-} c_i^\mu > \frac{1}{2} + \delta, \quad (16)$$

$$r^\mu = 0 \rightarrow c_j^\mu = \sum_{a \in A_\downarrow} \tilde{\kappa}_a \prod_{i \in \nu_a^-} c_i^\mu < \frac{1}{2} - \delta, \quad (17)$$

where $\tilde{\kappa}_a = \kappa_a / (\sum_{a \in A_\uparrow} \kappa_a)$.

Now consider a one-layer network of the R12 and R22 reactions with only two response variables ($N_R = 2$). The deterministic equations for the average concentrations of the two response variables in the thermodynamic limit $c_j^\mu, c_{j'}^\mu$ are

$$\frac{d}{dt}c_j^\mu = \frac{d}{dt}c_{j'}^\mu = \sum_{a \in A_\downarrow} \kappa_a \prod_{i \in \nu_a^-} c_i^\mu - \sum_{a \in A_\uparrow} \kappa_a c_j^\mu c_{j'}^\mu, \quad \forall \mu. \quad (18)$$




Figure 3. The response vs signal for the one-layer R11, R12, R21 and R22 reaction networks. The associated panels are: R11 = (a1), (b1), (c1), (d1), R12 = (a2), (b2), (c2), (d2), R21 = (a3), (b3), (c3), (d3), and R22 = (a4), (b4), (c4), (d4). The polarization p_R^μ and the response overlaps ($q_R^\mu, q_R^{\mu\mu'}$) are plotted in terms of the signals activity m_S^μ and the signals overlap $q_S^{\mu\mu'}$. The κ_a are chosen randomly and uniformly in $(0, 10)$ in fully connected networks of size $(N_S = 10, N_R = 5)$. The number of iterations is $t = 2000$ for the R11, R12, and $t = 4000$ for the R21, R22 networks. The points and colors show the values and frequencies of the observed data, respectively.

The stationary equations read

$$0 = \sum_{a \in A_\downarrow} \kappa_a \prod_{i \in \nu_a^-} c_i^\mu - \sum_{a \in A_\uparrow} \kappa_a (c_j^\mu)^2, \quad \forall \mu. \quad (19)$$

Here by symmetry the two responses have the same average concentrations. Then, by taking $(c_j^\mu)^2$ as the response, one obtains the same equations as equations (15)–(17) for the R11 and R21 networks. Therefore, in the following, we consider only the R11 and R21 reaction networks with a single response variable.

We take P random and independent signal-response patterns $\{\mathbf{s}^\mu, \mathbf{r}^\mu\}$. That is, for each $\mu = 1, \dots, P$ we set the $s_i^\mu = 0, 1$ and $r^\mu = 0, 1$ with equal probability. Then we use the least-squares method to check if there is a solution for the reaction rates which satisfy equations (15)–(17) for any μ . Figure 5 displays the probability of having a solution obtained by solving the above equations for at least 500 independent realizations of the patterns and the reaction networks. To have a fair comparison, besides the fully connected R11 and R21 networks, here we also report the results for a sparse R21 network with $M = N_S$ reversible reactions, where only a random subset of all the possible R21 reactions are present.

In figure 5 we observe that the number of stored patterns increases faster with N_S in the R21 reaction networks than the R11 network. More precisely, the number of stored patterns scales with N_S^ν , where the exponent $\nu \simeq 0.5$ in the R21 reaction networks which is nearly two times greater than that of the R11 network. The probability of having a solution of course increases with the number of signals in both the networks. But this probability is growing faster with N_S , and therefore with the number of reactions, in the R21 networks compared to the R11 network. This means that the two-signal interactions which are added by the R21 reactions are helpful in capturing some level of correlations in the signal patterns. It is well known that statistical models with two-body interactions are rich enough to capture strong correlations which are




Figure 4. The response vs signal for the two-layer R11, R12, R21 and R22 reaction networks. The associated panels are: R11 = (a1), (b1), (c1), (d1), R12 = (a2), (b2), (c2), (d2), R21 = (a3), (b3), (c3), (d3), and R22 = (a4), (b4), (c4), (d4). The polarization p_R^μ and the response overlaps ($q_R^\mu, q_R^{\mu\mu'}$) are plotted in terms of the signals activity m_S^μ and the signals overlap $q_S^{\mu\mu'}$. The κ_a are chosen randomly and uniformly in $(0, 10)$ in fully connected networks of size $(N_S = 10, N_h = 5, N_R = 5)$. The number of iterations is $t = 2000$ for the R11, R12, and $t = 4000$ for the R21, R22 networks. The points and colors show the values and frequencies of the observed data, respectively.

observed in complex systems and phase transitions. Note that these results are for deterministic reaction networks. To study the learning of stochastic reactions in a multilayer network, in the following we resort to a stochastic learning algorithm.

3.2. Learning the patterns


Consider P signal-response patterns $\{\mathbf{s}^\mu, \mathbf{r}^\mu\}$ and a stochastic reaction network with reaction rates $\kappa_a(0) \in (0, \kappa_{\max})$. We assume that the network structure is fixed. All the variables except the signal variables (in $\mathcal{S}$) are free to change in range $(0, X_{\max})$ by the system dynamics. The aim is to change gradually the network parameters to find the rates that store the above signal-response activities. We say the system learned pattern μ if the signal values s_i^μ result in the corresponding response concentrations that is, $|c_j^\mu - r_j^\mu| < \frac{1}{2} - \delta$ for all $j \in \mathcal{R}$. In practice, reaction rates can be altered by introducing catalysts, or changing catalytic capacity of an enzyme (via mutations or adding regulators) in a biochemical reaction. More precisely, this can be done by changing the concentration of an enzyme, or by changing the enzyme itself, as it happens during evolution (large time scales).

Given the reaction network's structure, we aim to identify the reaction rates that maximize the following objective function:

$$\mathcal{E}(\boldsymbol{\kappa}) = -\frac{1}{P} \sum_{\mu=1}^P E(\mathbf{c}^\mu : \mathbf{r}^\mu), \quad (20)$$

where $E(\mathbf{c}^\mu : \mathbf{r}^\mu)$ measures deviations from the expected responses $\mathbf{r}^\mu$ given the $\mathbf{s}^\mu$,

$$E(\mathbf{c}^\mu : \mathbf{r}^\mu) = \sum_{j \in \mathcal{R}} |c_j^\mu - r_j^\mu|. \quad (21)$$



By minimizing the above energy function for each pattern μ the system is pushed towards responses that are close to the expected ones for that pattern.

The idea is to update the reaction rates according to the sign of linear correlations between the products of a reaction $\sum_{i \in \nu_a^+} (2x_i^\mu - 1)$ and deviations in the responses $\sum_{j \in \mathcal{R}} (2r_j^\mu - 1)(2x_j^\mu - 1)$; a positive correlation means that increasing the products would result to a larger overlap between the observed responses and the expected ones.

We start with random initial parameters and run the system dynamics (i.e. Gillespie algorithm) for t_{eq} iterations to compute the objective function $\mathcal{E}(0)$. Then, for $n = 1, \dots, T$, we do the following for each pattern μ (randomly selected):

- compute linear correlations c_{ij}^μ from the statistics of pattern μ in time interval Δt
- for each reaction a compute $\Delta_a^\mu = \sum_{i \in \nu_a^+} \left(\sum_{j \in \mathcal{R}} (2r_j^\mu - 1) c_{ij}^\mu \right)$
- use Δ_a^μ to move the reaction rate κ_a from its current optimal value:
 - * with probability $1/2$ change $\kappa_a \rightarrow \kappa_a + \epsilon \Delta_a^\mu$, otherwise change it in a random direction
- run the dynamics for another t_{eq} iterations to determine the value of the objective function $\mathcal{E}(n)$ and accept the changes in the κ_a if the objective function shows an increase

We use the stochastic dynamics of section 2.3 and the above learning algorithm to store random and uncorrelated patterns of signals and responses. That is, the values $\{s_i^\mu = 0, 1 : i \in \mathcal{S}\}$ and $\{r_j^\mu = 0, 1 : j \in \mathcal{R}\}$ are chosen with equal probability and independently for $\mu = 1, \dots, P$. The signal patterns here are distinct. Recall that the response concentrations c_j^μ are determined by the reaction rates κ and the signal pattern $\mathbf{s}^\mu$. The aim of learning is to make the response concentrations c_j^μ as close as possible to the r_j^μ for all patterns μ . We checked that the above procedure works when we start from random reaction networks with known signal-response patterns; we obtain the responses for a number of input signals in a network with known reaction rates and then try to reconstruct the parameters by the above algorithm. Moreover, the algorithm performs better than a zero-temperature Monte Carlo which means that using the correlations to update the reaction rates is relevant here. For the parameters of the algorithm we take $t_{eq} = 2000-4000$, $\Delta t = 1000$, and $\epsilon = 0.02-0.2$. Figure 6 shows the average number of stored patterns in the one- and two-layer networks as the number of learning iterations increases. It is observed that the networks are able to learn in average more




Figure 6. Learning evolution of the fully connected one/two-layer R11 and R21 reaction networks. Top panels (a1)–(d1) and bottom panels (a2)–(d2) show the results for the one- and two-layer networks, respectively, for different numbers of patterns P and signals N_S . The relative number of stored patterns P_s/P is plotted vs the number of learning iterations T . The κ_{ai} are limited to $(0, 10)$ in fully connected networks of size $(N_S = 6, 8, 10, N_R = 1)$. The number of hidden variables in the two-layer networks is $N_h = N_S/2$.





Figure 7. Learning performance of the one-layer R11 and R21 reaction networks. The relative number of stored patterns P_s/P (after $T = 2000$ learning iterations) is plotted vs the number of patterns P . Panels (a) and (b) show the results for the fully connected networks. Panel (c) shows the results for a sparse R21 network with $M = N_S$ reversible reactions. The κ_{ai} are limited to $(0, 10)$ in networks of size $(N_S = 6, 8, 10, N_R = 1)$.

than half of the presented patterns (P) after a few tens of iterations. However, except for the case of a single pattern ($P = 1$), the learning is far from perfect with a large gap separating the cases $P \geq 2$. This is the case for both the R11 and R21 networks as shown in figure 6. Note that here we are working with random signal-response patterns. From the study of random reaction networks in section 3.1 we know that it would be easier to store patterns which are closer to the expected functions of the networks. For instance, patterns in which the activity of responses is low are better represented by the R21 reaction networks.

Figure 7 displays the average number of stored patterns $\langle P_s \rangle$ after $T = 2000$ learning iterations as a function of P . We see that $\langle P_s \rangle$ approaches the value expected from a random network after a few number of patterns. As before, the R21 reaction network exhibits a better storage performance than the R11 networks. That is, the chance of storing more than half of the presented patterns decreases more rapidly with P in the R11 network. This is the case also in the sparse R21 network with the same number of reactions as in the R11 network.

3.3. Robustness

Chemical reaction systems like cellular signaling networks may undergo aging or disease processes which comprise of gradual accumulation of defects. A defect could for instance alter the concentration of an enzyme, which indirectly controls the reaction rate of a biochemical transformation, or even modify the structure of the reaction network by introducing new reactions to the system. Here we consider only the




former defects, which are simply represented by small changes in the reaction rates of a network. In this section, we want to see how much the stored patterns in the above stochastic reaction networks are robust against random variations in the reaction rates.

To model such a process, we begin with the stationary state of an optimal (so-called ‘healthy’) reaction network, where a random signal-response pattern is stored (i.e. $P = 1$). The value of the associated objective function is denoted by $\mathcal{E}_{\text{old}}$. The time evolution of a disease then is modeled in the following way [49]: For $t_d = 1, \dots, t_{\text{max}}$ do:

- each reaction rate changes with (mutation) probability $\alpha(t_d)$ from κ_a to $\kappa_a \pm \delta\kappa$
- compute the new objective function $\mathcal{E}_{\text{new}}$
- accept the changes in the reaction rates if the objective function increases, otherwise, accept the changes with probability $\exp(\beta(t)\Delta\mathcal{E})$. If the changes are accepted set $\mathcal{E}_{\text{old}} = \mathcal{E}_{\text{new}}$

The mutation probability $\alpha(t_d)$ is expected to increase with the disease time step t_d . On the other hand, the parameter $\beta(t_d)$ which controls the acceptance probability is expected to decrease with t_d . Here $\beta(t_d) \geq 0$ plays the role of an inverse temperature in a reverse simulated annealing algorithm [51]. While the probability $\alpha(t_d)$ controls the rate of local changes (‘mutations’) in the reaction network, the global parameter $\beta(t_d)$ determines the system susceptibility to such local mutations (‘immunity’). The direct



simulated annealing algorithm, where $\alpha(t_d)$ decreases and $\beta(t_d)$ increases with time, can also be employed as a learning algorithm.

Note that in each time step t_d all the reactions have the chance to increase or decrease with equal probability by $\delta\kappa$. In the following numerical simulations, we set $\delta\kappa = 0.5$, and as before we limit the reactions rates $\kappa_a \in (0, 10)$. Every time that a change in the reaction rates is suggested, we run the Gillespie algorithm for $t_{eq} = 2000 - 4000$ iterations to reach the stationary state, where the new objective function is estimated.

In figure 8, we check the robustness of the stored patterns against the disease evolution in the one-layer networks. The learning algorithm of section 3.2 was used to store a single random pattern in $T = 2000$ learning iterations. We see that the survival time of the stored pattern is longer in the R11 network than the R21 network for different rates of disease evolution. Thus, the basin of attraction of the stored pattern should be larger in the R11 network. Moreover, it is observed that the transition from the healthy state to the disease state gets sharper by increasing the size of the R11 network. Whereas in the R21 network the transition becomes smoother as the number of variables grows. This shows that the stored pattern in the R11 network is well separated from other spurious states in the reaction rate space. Similar results for the two-layer networks are reported in figure 9. In addition to the above observations, here the system displays a two-stage behavior when the time scales of the $\alpha(t_d)$ and $\beta(t_d)$ are different. More precisely, the change in the mutation probability in a two-layer network has a considerable effect on the system even for large values of β . This

means that the space of reaction rates around the stored pattern in the one-layer and two-layer networks are different. The difference could be due to the fact that learning of two-layer networks with the R21 reactions is more difficult than that of the one-layer networks.

4. Conclusion and discussion

In this paper we studied the function of multilayer networks of reversible and stochastic reactions connecting the subsets of signal, response, and hidden species. We used elementary reactions R_{mn} connecting $m = 1, 2$ reactants to $n = 1, 2$ products and employed the Gillespie algorithm to simulate the stochastic dynamics of the species. We observed that the typical response function of the R11 and R22 networks with uniformly random reaction rates is concentrated around the activity $c = 0.5$. On the other hand, the R12 and R21 reaction networks displayed high and low response activities, respectively, with greater response polarization and larger overlap between the responses compared to the R11 and R22 networks. This effect was enhanced in the two-layer networks of the same reactions. Moreover, we found that the average polarization and similarity of responses in the R12 (R21) network increase (decrease) with the number of active signals, in contrast to the symmetric and non-monotonic behavior of the R11 and R22 networks. In summary, it seems that the function space of the random reaction networks is represented by simple functions very much like random neural networks [16]. To observe more complex functions, it would be essential to introduce recurrent reactions connecting reactant and product species in hidden layer of the networks.

We found that the average number of stored patterns in deterministic reaction networks grows faster with the number of signals in the R21 reaction networks than in the R11 networks. This highlights the role of two-signal interaction that are induced by the R21 reactions. It remains to be seen how much this effect is changed by increasing the number of hidden layers. We also introduced a stochastic learning algorithm to store a number of random signal-response activities in the stochastic reaction networks. The algorithm was based on the correlations between the products of a reaction and the responses. Here, we found it very difficult to store an extensive number of random signal-response patterns in both the one- and two-layer networks, at least with the above learning algorithm. Still, the R21 reaction network performed better than the R11 network in learning the random patterns. The patterns in the R11 networks are however more robust with respect to random changes in the reaction rates in a stochastic process of defect evolution. Moreover, the two-layer R21 reaction networks are more sensitive to changes in the mutation probability perhaps due to the appearance of other metastable states in the learning process.

In this study, we considered elementary reactions which in contrast to interactions in artificial neural networks are always excitatory in the sense that more reactants result to more products. It is of course possible to model effectively the inhibitory interactions by introducing other auxiliary species to the system [34]. Adding this feature to the reactions will expectedly enhance the learning performance of these networks at the expense of working with a larger number of species. In addition, the stochastic algorithm that we used to store the signal-response patterns can be improved to obtain better lower bounds for the learning capacity of artificial reaction networks. For instance, one could employ more efficient algorithms than the Gillespie algorithm for simulation of the system dynamics [43–45].

Please note that the reactions in this work are stochastic, which sets this study apart from other studies on learning deterministic reaction networks [29, 30, 34, 39]. The stochastic nature of the reactions increases the computational complexity of the problem but also provides more information about the probability distribution of the species. We used this information for example in the learning algorithm of section 3.2 to update the reaction rates according to the values of the correlations between the products of a reaction and the response species. Information about the statistical dependencies of the species, which is absent in deterministic models, can be useful in a learning algorithm to devise more elaborate strategies to update the reaction rates. In [appendix](#) we use a similar strategy to store random patterns of signals and responses in an enzymatic reaction network.

The purpose of this study was to provide a theoretical framework for future applications, such as diagnostic problems. For instance, suppose we have access to the history of concentration variations of

species in a reaction network that is developing an undesired state, so-called a ‘disease’ state. The starting reaction network could be in a ‘healthy’ state that stores several signal-response patterns. To diagnose the developing disease from such dynamic data, one must first understand how the initial system of interacting species represents the stored patterns, as well as how robust these healthy states are against relevant mutations in the reaction network. A signal transduction reaction network, like the one studied here, enables us to investigate disease dynamics and the effectiveness of diagnostic algorithms in a simple benchmark model.

Data availability statement

All data that support the findings of this study are included within the article (and any supplementary files).

Acknowledgment

This work was performed using the Academic Leiden Interdisciplinary Cluster Environment (ALICE), the compute resources provided by Leiden University.

Appendix. Learning of a stochastic enzymatic reaction network

In this section we provide an example of learning by changing the concentrations of enzymes in a stochastic enzymatic reaction network. We consider reactions that connect a substrate to a product through a complex formation





Figure 10 schematically displays a reaction network of N_S substrates (signals) connected by the same number of distinct enzymes and complexes to a single product (response). Here, in contrast to the main text, we try to store random patterns of signals and responses by changing only the concentrations of the enzymes. As before the kinetics is given by the mass action law but with fixed reaction rates κ_a which are distributed randomly and uniformly in $(0, 10)$. The results at the end are averaged over independent realizations of the reaction rates and the stochastic dynamics.

The learning algorithm is the same as the one presented in section 3.2, except that here we use the correlations between the enzymes and the product to change the enzyme concentrations. We begin with random initial numbers for the enzymes and execute the Gillespie algorithm for t_{eq} iterations to compute the objective function $\mathcal{E}(0)$ (from equations (20) and (21)). Then, for $n = 1, \dots, T$, we do the following for each pattern μ :

- compute linear correlations c_{ij}^μ from the statistics of pattern μ in time interval Δt
- for each enzyme i compute $\Delta_i^\mu = \sum_{j \in \mathcal{R}} (2r_j^\mu - 1) c_{ij}^\mu$
- use Δ_i^μ to move the number of enzymes X_i from its current optimal value:
 - * with probability 1/2 change $X_i \rightarrow X_i + \epsilon \Delta_i^\mu$, otherwise change it in a random direction
- run the dynamics for another t_{eq} iterations to determine the value of the objective function $\mathcal{E}(n)$ and accept the changes in the X_i if the objective function shows an increase

Figure 11 shows the average number of stored patterns P_s when we try to store P random patterns in this way. The parameters used in the algorithm are $t_{eq} = 3000$, $\Delta t = 1000$, and $\epsilon = 0.02X_{\max}$. Note that the reaction rates are fixed to random values and we do not tune the reaction rates. Therefore, learning is more difficult here compared to the case we studied in the main text.



ORCID iDs

Abolfazl Ramezanpour <https://orcid.org/0000-0002-3650-7722>

Alireza Mashaghi <https://orcid.org/0000-0002-2157-1211>

References

- [1] Stocker S, Csányi G, Reuter K and Margraf J T 2020 Machine learning in chemical reaction space *Nat. Commun.* **11** 5505
- [2] Evers T M J, Holt L J, Alberti S and Mashaghi A 2021 Reciprocal regulation of cellular mechanics and metabolism *Nat. Metab.* **3** 456–68
- [3] Araujo R P and Liotta L A 2018 The topological requirements for robust perfect adaptation in networks of any size *Nat. Commun.* **9** 1757
- [4] Reed M, Best J, Golubitsky M, Stewart I and Nijhout H F 2017 Analysis of homeostatic mechanisms in biochemical networks *Bull. Math. Biol.* **79** 2534–57
- [5] Aoki S K, Lillacci G, Gupta A, Baumschlager A, Schweingruber D and Khammash M 2019 A universal biomolecular integral feedback controller for robust perfect adaptation *Nature* **570** 533–7
- [6] Wong A S Y, Pogodaev A A, Vialshin I N, Helwig B and Huck W T S 2017 Molecular engineering of robustness and resilience in enzymatic reaction networks *J. Am. Chem. Soc.* **139** 8146–51
- [7] Hopfield J J 1982 Neural networks and physical systems with emergent collective computational abilities *Proc. Natl Acad. Sci.* **79** 2554–8

- [8] LeCun Y, Touresky D, Hinton G and Sejnowski T 1988 A theoretical framework for back-propagation *Proc. 1988 Connectionist Models Summer School* vol 1 pp 21–28
- [9] Lillicrap T P, Santoro A, Marris L, Akerman C J and Hinton G 2020 Backpropagation and the brain *Nat. Rev. Neurosci.* **21** 335–46
- [10] Barkai E, Hansel D and Kanter I 1990 Statistical mechanics of a multilayered neural network *Phys. Rev. Lett.* **65** 2312
- [11] Barkai E, Hansel D and Sompolinsky H 1992 Broken symmetries in multilayered perceptrons *Phys. Rev. A* **45** 4146
- [12] Monasson R and Zecchina R 1995 Weight space structure and internal representations: a direct approach to learning and generalization in multilayer neural networks *Phys. Rev. Lett.* **75** 2432
- [13] Braunstein A and Zecchina R 2006 Learning by message passing in networks of discrete synapses *Phys. Rev. Lett.* **96** 030201
- [14] Li B and Saad D 2018 Exploring the function space of deep-learning machines *Phys. Rev. Lett.* **120** 248301
- [15] Li B and Saad D 2020 Large deviation analysis of function sensitivity in random deep neural networks *J. Phys. A: Math. Theor.* **53** 104002
- [16] De Palma G, Kiani B and Lloyd S 2019 Random deep neural networks are biased towards simple functions *Advances in Neural Information Processing Systems* vol 32
- [17] Aubin B, Perkins W and Zdeborova L 2019 Storage capacity in symmetric binary perceptrons *J. Phys. A: Math. Theor.* **52** 294003
- [18] Borra F, Lagomarsino M C, Rotondo P and Gherardi M 2019 Generalization from correlated sets of patterns in the perceptron *J. Phys. A: Math. Theor.* **52** 384004
- [19] Rotondo P, Pastore M and Gherardi M 2020 Beyond the storage capacity: data-driven satisfiability transition *Phys. Rev. Lett.* **125** 120601
- [20] Magnasco M O 1997 Chemical kinetics is Turing universal *Phys. Rev. Lett.* **78** 1190
- [21] Soloveichik D, Cook M, Winfree E and Bruck J 2008 Computation with finite stochastic chemical reaction networks *Nat. Comput.* **7** 615–33
- [22] Cook M, Soloveichik D, Winfree E and Bruck J 2009 Programmability of chemical reaction networks *Algorithmic Bioprocesses* (Springer) pp 543–84
- [23] Morgan C, Stefanovic D, Moore C and Stojanovic M N 2004 Building the components for a biomolecular computer *Int. Workshop on DNA-Based Computers* (Springer) pp 247–57
- [24] Soloveichik D, Seelig G and Winfree E 2010 DNA as a universal substrate for chemical kinetics *Proc. Natl Acad. Sci.* **107** 5393–8
- [25] Yahiro W, Aubert-Kato N and Hagiya M 2018 A reservoir computing approach for molecular computing *Artificial Life Conf. Proc.* (MIT Press) pp 31–38
- [26] Nguyen H, Banda P, Stefanovic D and Teuscher C 2020 Reservoir computing with random chemical systems *Artificial Life Conf. Proc.* (MIT Press) pp 491–9
- [27] Bray D and Lay S 1994 Computer simulated evolution of a network of cell-signaling molecules *Biophys. J.* **66** 972–7
- [28] Duarte N C, Becker S A, Jamshidi N, Thiele I, Mo M L, Vo T D, Srivas R and BØ P 2007 Global reconstruction of the human metabolic network based on genomic and bibliomic data *Proc. Natl Acad. Sci.* **104** 1777–82
- [29] McGregor S, Vasas V, Husbands P and Fernando C 2012 Evolution of associative learning in chemical networks *PLoS Comput. Biol.* **8** e1002739
- [30] Banda P, Teuscher C and Lakin M R 2013 Online learning in a chemical perceptron *Artif. Life* **19** 195–219
- [31] Lakin M R, Minnich A, Lane T and Stefanovic D 2014 Design of a biochemical circuit motif for learning linear functions *J. R. Soc. Interface* **11** 20140902
- [32] Lakin M R and Stefanovic D 2015 Supervised learning in an adaptive DNA strand displacement circuit *Int. Workshop on DNA-Based Computers* (Springer) pp 154–67
- [33] Poole W, Ortiz-Munoz A, Behera A, Jones N S, Ouldrige T E, Winfree E and Gopalkrishnan M 2017 Chemical boltzmann machines *Int. Conf. on DNA-Based Computers* (Springer) pp 210–31
- [34] Blount D, Banda P, Teuscher C and Stefanovic D 2017 Feedforward chemical neural network: an in silico chemical system that learns XOR *Artif. Life* **23** 295–317
- [35] Braunstein A, Muntioni A P and Pagnani A 2017 An analytic approximation of the feasible space of metabolic networks *Nat. Commun.* **8** 1–9
- [36] Wong A S and Huck W T 2017 Grip on complexity in chemical reaction networks *Beilstein J. Org. Chem.* **13** 1486–97
- [37] Moorman A, Samaniego C C, Maley C and Weiss R 2019 A dynamical biomolecular neural network 2019 *IEEE 58th Conf. on Decision and Control (CDC)* (IEEE) pp 1797–802
- [38] Vasic M, Chalk C, Khurshid S and Soloveichik D 2020 Deep molecular programming: a natural implementation of binary-weight ReLU neural networks *Int. Conf. on Machine Learning* (PMLR) pp 9701–11
- [39] Anderson D F, Joshi B and Deshpande A 2021 On reaction network implementations of neural networks *J. R. Soc. Interface* **18** 20210031
- [40] Feinberg M 2019 *Foundations of Chemical Reaction Network Theory* vol 10 (Springer International Publishing) pp 978–3
- [41] Horowitz J M and England J L 2017 Spontaneous fine-tuning to environment in many-species chemical reaction networks *Proc. Natl Acad. Sci.* **114** 7565–70
- [42] Gillespie D T 2007 Stochastic simulation of chemical kinetics *Annu. Rev. Phys. Chem.* **58** 35–55
- [43] Ramaswamy R, González-Segredo N and Szbalzarini I F 2009 A new class of highly efficient exact stochastic simulation algorithms for chemical reaction networks *J. Chem. Phys.* **130** 244104
- [44] Gupta A, Briat C and Khammash M 2014 A scalable computational framework for establishing long-term behavior of stochastic reaction networks *PLoS Comput. Biol.* **10** e1003669
- [45] Hepp B, Gupta A and Khammash M 2015 Adaptive hybrid simulations for multiscale stochastic reaction networks *J. Chem. Phys.* **142** 034118
- [46] Shinar G and Feinberg M 2010 Structural sources of robustness in biochemical reaction networks *Science* **327** 1389–91
- [47] Donzé A, Fanchon E, Gattepaille L M, Maler O and Tracqui P 2011 Robustness analysis and behavior discrimination in enzymatic reaction networks *PLoS One* **6** e24246
- [48] Anderson D F, Enciso G A and Johnston M D 2014 Stochastic analysis of biochemical reaction networks with absolute concentration robustness *J. R. Soc. Interface* **11** 20130943
- [49] Ramezanpour A and Mashaghi A 2020 Disease evolution in reaction networks: implications for a diagnostic problem *PLoS Comput. Biol.* **16** e1007889
- [50] Goutsias J and Jenkinson G 2013 Markovian dynamics on complex reaction networks *Phys. Rep.* **529** 199–264
- [51] Kirkpatrick S, Gelatt C D and Vecchi M P 1983 Optimization by simulated annealing *Science* **220** 671–80