



OPEN ACCESS

EDITED BY

Byungjoon Min,
Chungbuk National University, Republic of
Korea

REVIEWED BY

Yukio Pegio Gunji,
Waseda University, Japan
Valentina Lanza,
University of Le Havre, France

*CORRESPONDENCE

Patrick Krauss,
✉ patrick.krauss@fau.de

RECEIVED 12 August 2024

ACCEPTED 20 September 2024

PUBLISHED 08 October 2024

CITATION

Metzner C, Schilling A, Maier A and Krauss P
(2024) Recurrence resonance - noise-
enhanced dynamics in recurrent
neural networks.
Front. Complex Syst. 2:1479417.
doi: 10.3389/fcpxs.2024.1479417

COPYRIGHT

© 2024 Metzner, Schilling, Maier and Krauss.
This is an open-access article distributed under
the terms of the [Creative Commons Attribution
License \(CC BY\)](#). The use, distribution or
reproduction in other forums is permitted,
provided the original author(s) and the
copyright owner(s) are credited and that the
original publication in this journal is cited, in
accordance with accepted academic practice.
No use, distribution or reproduction is
permitted which does not comply with these
terms.

Recurrence resonance - noise-enhanced dynamics in recurrent neural networks

Claus Metzner¹, Achim Schilling^{1,2}, Andreas Maier¹ and
Patrick Krauss^{1,2*}

¹Cognitive Computational Neuroscience Group, Pattern Recognition Lab, Friedrich-Alexander-
University Erlangen-Nürnberg (FAU), Erlangen, Germany, ²Neuroscience Lab, University Hospital
Erlangen, Erlangen, Germany

Understanding how neural networks process information is a fundamental challenge in neuroscience and artificial intelligence. A pivotal question in this context is how external stimuli, particularly noise, influence the dynamics and information flow within these networks. Traditionally, noise is perceived as a hindrance to information processing, introducing randomness and diminishing the fidelity of neural signals. However, distinguishing noise from structured input uncovers a paradoxical insight: under specific conditions, noise can actually enhance information processing. This intriguing possibility prompts a deeper investigation into the nuanced role of noise within neural networks. In specific motifs of three recurrently connected neurons with probabilistic response, the spontaneous information flux, defined as the mutual information between subsequent states, has been shown to increase by adding ongoing white noise of some optimal strength to each of the neurons. However, the precise conditions for and mechanisms of this phenomenon called 'recurrence resonance' (RR) remain largely unexplored. Using Boltzmann machines of different sizes and with various types of weight matrices, we show that RR can generally occur when a system has multiple dynamical attractors, but is trapped in one or a few of them. In probabilistic networks, the phenomenon is bound to a suitable observation time scale, as the system could autonomously access its entire attractor landscape even without the help of external noise, given enough time. Yet, even in large systems, where time scales for observing RR in the full network become too long, the resonance can still be detected in small subsets of neurons. Finally, we show that short noise pulses can be used to transfer recurrent neural networks, both probabilistic and deterministic, between their dynamical attractors. Our results are relevant to the fields of reservoir computing and neuroscience, where controlled noise may turn out a key factor for efficient information processing leading to more robust and adaptable systems.

KEYWORDS

recurrent neural networks (RNN), resonance, information processing, dynamics, noise

1 Introduction

Artificial neural networks are a cornerstone of many contemporary machine learning methods, especially in deep learning (LeCun et al., 2015). Over the past decades, these systems have found extensive applications in both industrial and scientific domains (Alzubaidi et al., 2021). Typically, neural networks in machine learning are organized

in layered structures, where information flows unidirectionally from the input layer to the output layer. In contrast, Recurrent Neural Networks (RNNs) incorporate feedback loops within their neuronal connections, allowing information to continuously circulate within the system (Maheswaranathan et al., 2019). Consequently, RNNs function as autonomous dynamical systems with ongoing neural activity even in the absence of external input, and they are recognized as ‘universal approximators’ (Maximilian Schäfer and Zimmermann, 2006). These unique characteristics have spurred a significant increase in research on artificial RNNs, leading to both advancements and intriguing unresolved issues: Thanks to their recurrent connectivity, RNNs are particularly well-suited for processing time series data (Jaeger, 2001) and for storing sequential inputs over time (Schuecker et al., 2018; Büsing et al., 2010; Dambre et al., 2012; Wallace et al., 2013; Gonon and Ortega, 2021). For example, RNNs have been shown to learn robust representations by dynamically balancing compression and expansion (Farrell et al., 2022). Specifically, a dynamic state known as the ‘edge of chaos’, situated at the transition between periodic and chaotic behavior (Kadmon and Sompolinsky, 2015), has been extensively investigated and identified as crucial for computation (Wang et al., 2011; Boedecker et al., 2012; Langton, 1990; Natschläger et al., 2005; Legenstein and Maass, 2007; Bertschinger and Natschläger, 2004; Schrauwen et al., 2009; Toyozumi and Abbott, 2011; Kaneko and Suzuki, 1994; Solé and Miramontes, 1995) and short-term memory (Haruna and Nakajima, 2019; Ichikawa and Kaneko, 2021). Moreover, several studies focus on controlling the dynamics of RNNs (Rajan et al., 2010; Jaeger, 2014; Haviv et al., 2019), particularly through the influence of external or internal noise (Molgedey et al., 1992; Ikemoto et al., 2018; Krauss et al., 2019a; Bönsel et al., 2022; Metzner and Krauss, 2022). RNNs are also proposed as versatile tools in neuroscience research (Barak, 2017). Notably, very sparse RNNs, similar to those found in the human brain (Song et al., 2005), exhibit remarkable properties such as superior information storage capacities (Brunel, 2016) (Narang et al., 2017; Gerum et al., 2020; Folli et al., 2018).

In our previous research, we systematically analyzed the relation between network structure and dynamical properties in recurrent three-neuron motifs (Krauss et al., 2019b). We also demonstrated how statistical parameters of the weight matrix can be used to control the dynamics in large RNNs (Krauss et al., 2019c; Metzner and Krauss, 2022). Another focus of our research are noise-induced resonance phenomena (Bönsel et al., 2022; Schilling et al., 2022; Krauss et al., 2016; Schilling et al., 2021; Schilling et al., 2023). In particular, we discovered that in specific recurrent motifs of three probabilistic neurons, connected with ternary $(-1, 0, +1)$ weights, the mutual information I between subsequent system states can go through a resonance-like maximum when normal-distributed white noise of increasing standard deviation r is added independently to all neurons. The phenomenon was called ‘Recurrence Resonance’ (RR) (Krauss et al., 2019a), because I can be interpreted as the spontaneous recurrent information flux in the network. It grows with the number of visited system states and with the degree of predictability of each successor state from its predecessor.

Since I is a key factor for the information processing faculties of RNNs, it became important to understand which types of weight matrices enable a large spontaneous information flux in probabilistic RNNs, such as Symmetric Boltzmann Machines (SBMs, see

Methods for details). By reverse-engineering evolutionary optimized networks (Gerum et al., 2020), we found indeed a universal design principle for maximizing I (Metzner et al., 2024). It was called ‘NRooks’, because in optimal N -neuron networks each row and column of the $N \times N$ weight matrix only contains a single non-zero entry, resembling the N -rooks-problem in chess (Katz and Sobel, 1972). While these N non-zero elements should ideally have the same magnitude w , their signs can be arbitrary. In the limit of large magnitudes w , the SBMs become quasi-deterministic and the information flux I approaches the theoretical maximum $I_{max} = N$. In this extreme case, all 2^N possible system states are periodically visited in a fixed order.

The present work aims to understand, on a deeper level than before, the pre-conditions of the RR phenomenon, as well as its mechanism. As model systems, we will mainly use probabilistic SBMs, but we will also briefly consider deterministic networks with ‘tanh’ activation functions. Different types of weight matrices will be investigated, but NRooks system will play a particularly important role, as their information theoretic properties are very well understood (Metzner et al., 2024).

2 Methods

2.1 Neural network model

We consider a recurrent network of N_{neu} model neurons. The total sum of inputs entering neuron n in the discrete time step t is given by

$$u_n^{(t)} = \left(v_n + \sum_{m=1}^{N_{neu}} w_{nm} s_m^{(t)} \right) + (q \ x_n^{(t)}) + (r \ \eta_n^{(t)}). \quad (1)$$

Here, the first bracket contains a possible bias v_n , as well as a weighted sum of the momentary output signals $s_m^{(t)}$ from all neurons m in the network. The weighting w_{nm} describes the coupling strength from source neuron m to target neuron n . The second bracket accounts for the momentary external input signal $x_n^{(t)}$ entering neuron n , scaled by a global input coupling parameter q . Finally, the third bracket accounts for a random signal $\eta_n^{(t)}$ entering neuron n , scaled by a global **noise strength** parameter r . The $\eta_n^{(t)} \sim \mathcal{N}(0, 1)$ are statistically independent random numbers, drawn from a standard Gaussian distribution with mean 0 and standard deviation 1. Further on, we will denote the complete **weight matrix** by \mathbf{W} , the bias vector by \mathbf{v} , the momentary state vector by $\mathbf{s}^{(t)}$, and the external input vector by $\mathbf{x}^{(t)}$. The total sum $u_n^{(t)}$ of inputs, as defined in Equation 1, is used to compute the next output state $s_n^{(t+1)}$ of neuron n . This **update**, which is **performed synchronously** for all neurons, differs for the two models considered in this paper:

In the **deterministic model**, neural output signals are continuous in the range $[-1, +1]$ and are computed directly as the hyperbolic tangent of the total sum:

$$s_n^{(t+1)} = \tanh(u_n^{(t)}). \quad (2)$$

To initialize the deterministic network, the N_{neu} elements of the zero-time state vector $\mathbf{s}^{t=0} = (s_0^0, s_1^0, \dots, s_{N_{neu}-1}^0)$ are drawn independently from a uniform distribution in the range $[-1, +1]$.

In the **probabilistic model**, neural output signals are discrete with the two possible values $\{-1, +1\}$. The probability for the value $+1$, also called the ‘on’-probability, is computed as a logistic function of the total sum:

$$P_{on,n}^{(t+1)} = \sigma(u_n^{(t)}) = \frac{1}{1 + \exp(-u_n^{(t)})}. \quad (3)$$

To initialize the probabilistic network, the N_{neu} elements of the zero-time state vector $\mathbf{s}^{t=0} = (s_0^0, s_1^0, \dots, s_{N_{neu}-1}^0)$ are drawn independently from a Bernoulli distribution, in which the possible outcomes -1 and $+1$ occur with equal probability.

We also refer to our probabilistic model as a **Symmetrical Boltzmann Machine (SBM)**, which is called ‘symmetric’ because the binary outputs are set to $\{-1, +1\}$, rather than the values $\{0, 1\}$ in conventional Boltzmann machines. Our choice makes the SBMs directly comparable to deterministic tanh-networks with the same weight matrix.

Note that here we do not apply any input to the recurrent neural networks, and thus $q = 0$. Also, we do not use any biases, so that $v_n = 0$ as well. For some types of sparse weight matrices considered in this work, the elements have ternary values $w_{nm} \in \{-w, 0, +w\}$. In this case, w is called the **weight magnitude** parameter. The same name is also used for a multiplicative parameter w that scales the standard deviation of a weight matrix with originally random normal elements $w_{nm} \sim \mathcal{N}(0, 1)$.

After defining the weight matrix and randomly initializing a network, the time series of global system states $\mathbf{s}^{t=0}, \mathbf{s}^{t=1}, \dots$ is computed numerically for N_T time steps, a parameter also called the observation **time scale**.

2.2 Information theoretic quantities

Numerical evaluation of information theoretic quantities requires data with discrete values. The binary output of the SBM is perfectly suited for this purpose, but in the case of the tanh-network we first needed to binarize the continuous outputs $s_n^{(t)} \in [-1, +1] \rightarrow x_n^{(t)} \in \{0, 1\}$, defining $x = 0$ if $s < 0$, and $x = 1$ if $s \geq 0$. Thus, for both types of networks, the output time series can be reduced to a binary matrix $\mathbf{X} \cong \mathbf{s}^{(t)}$, or $\mathbf{Y} \cong \mathbf{s}^{(t+1)}$ for the time-shifted series. The matrix \mathbf{X} has N_T rows, each with N_{neu} binary values 0 or 1. The rows correspond to momentary global states x of the network, and (due to the binarization) the total number of possible global states is $N_x = 2^{N_{neu}}$.

The starting point for all our information theoretic quantities is the **joint probability** $P(x, y)$ that a global system state x is followed by a subsequent state y in the time series. Since the size N_x of the state space increases exponentially with the network size N_{neu} , the estimation of $P(x, y)$ becomes quickly unfeasible for large systems: Not only does it take too long before the system has ergodically spread over its entire state space and tried out all possible transitions $x \rightarrow y$, but in large systems it even becomes difficult to hold the huge matrix $P(x, y)$ in the computer memory. To alleviate the memory problem, we hold only the matrix elements between the subset of global system states that have actually been visited by the system within the time scale N_T , which can be much smaller if the system is trapped in a dynamical attractor. From the joint

probability $P(x, y)$ we directly obtain the two **marginal probabilities** $P(x)$ and $P(y)$, which in our case are practically identical, because each final state becomes an initial state in the next time step.

The first information theoretical quantity of interest in the **state entropy** $H(X)$ of global system states, defined by

$$H(X) = H(Y) = - \sum_x P(x) \log_2 P(x), \quad (4)$$

where all terms with $P(x) = 0$ count as zero.

The next relevant quantity is the **mutual information** $I(X; Y)$ between subsequent system states, defined by

$$I(X; Y) = \sum_x \sum_y P(x, y) \log_2 \left(\frac{P(x, y)}{P(x)P(y)} \right), \quad (5)$$

where all terms with $P(x) = 0$, or $P(y) = 0$, or $P(x, y) = 0$ count as zero.

The final important quantity is the conditional entropy $H(Y|X) = H(Y) - I(X; Y)$, which in our case can also be written as $H(Y|X) = H(X) - I(X; Y)$. This conditional entropy $H(Y|X)$ describes the random divergence from a specific initial state x to several possible final states y and is therefore called the **state divergence** $D = H - I$ in the following. A value of $D = 0$ would indicate perfectly deterministic behavior, a value of $D = H$ perfectly random behavior.

3 Results

3.1 Conditions and mechanisms of recurrence resonance

Our first goal is to identify the preconditions of RR, in particular regarding the network weight matrices. For this purpose, we consider a Symmetric Boltzmann Machine (SBM, see Methods for details) (Equations 1–3) with $N_{neu} = 5$ binary neurons. Since such a system has only $N_x = 2^5 = 32$ global states x , the entropy $H(X)$, the mutual information $I(X; Y)$ between subsequent system states x and y , as well as the divergence $D = H - I$ can be accurately estimated from the network’s output time series on a manageable time scale of $N_T = 10^4$ time steps (See Methods for details) (Equations 4, 5).

As a basis for the network’s weight matrix \mathbf{W} , the 25 elements are first drawn independently from a standard normal distribution $\sim \mathcal{N}(0, 1)$ (Figure 1A). This ‘frozen random matrix’ is then multiplied (scaled) with a weight magnitude parameter w , thus tuning the standard deviation of the matrix, while keeping its fundamental structure invariant. We can then explore how the three information theoretic quantities H , I and D depend on the weight magnitude w and on the strength r of added white noise (See Methods for details).

Before that, it is useful to imagine the dynamical structure of the SBM as a state transition graph, in which the 32 nodes represent the possible global states

$$x \in \{(0, 0, 0, 0, 0) \equiv 0, (0, 0, 0, 0, 1) \equiv 1, \dots, (1, 1, 1, 1, 1) \equiv 31\}, \quad (6)$$

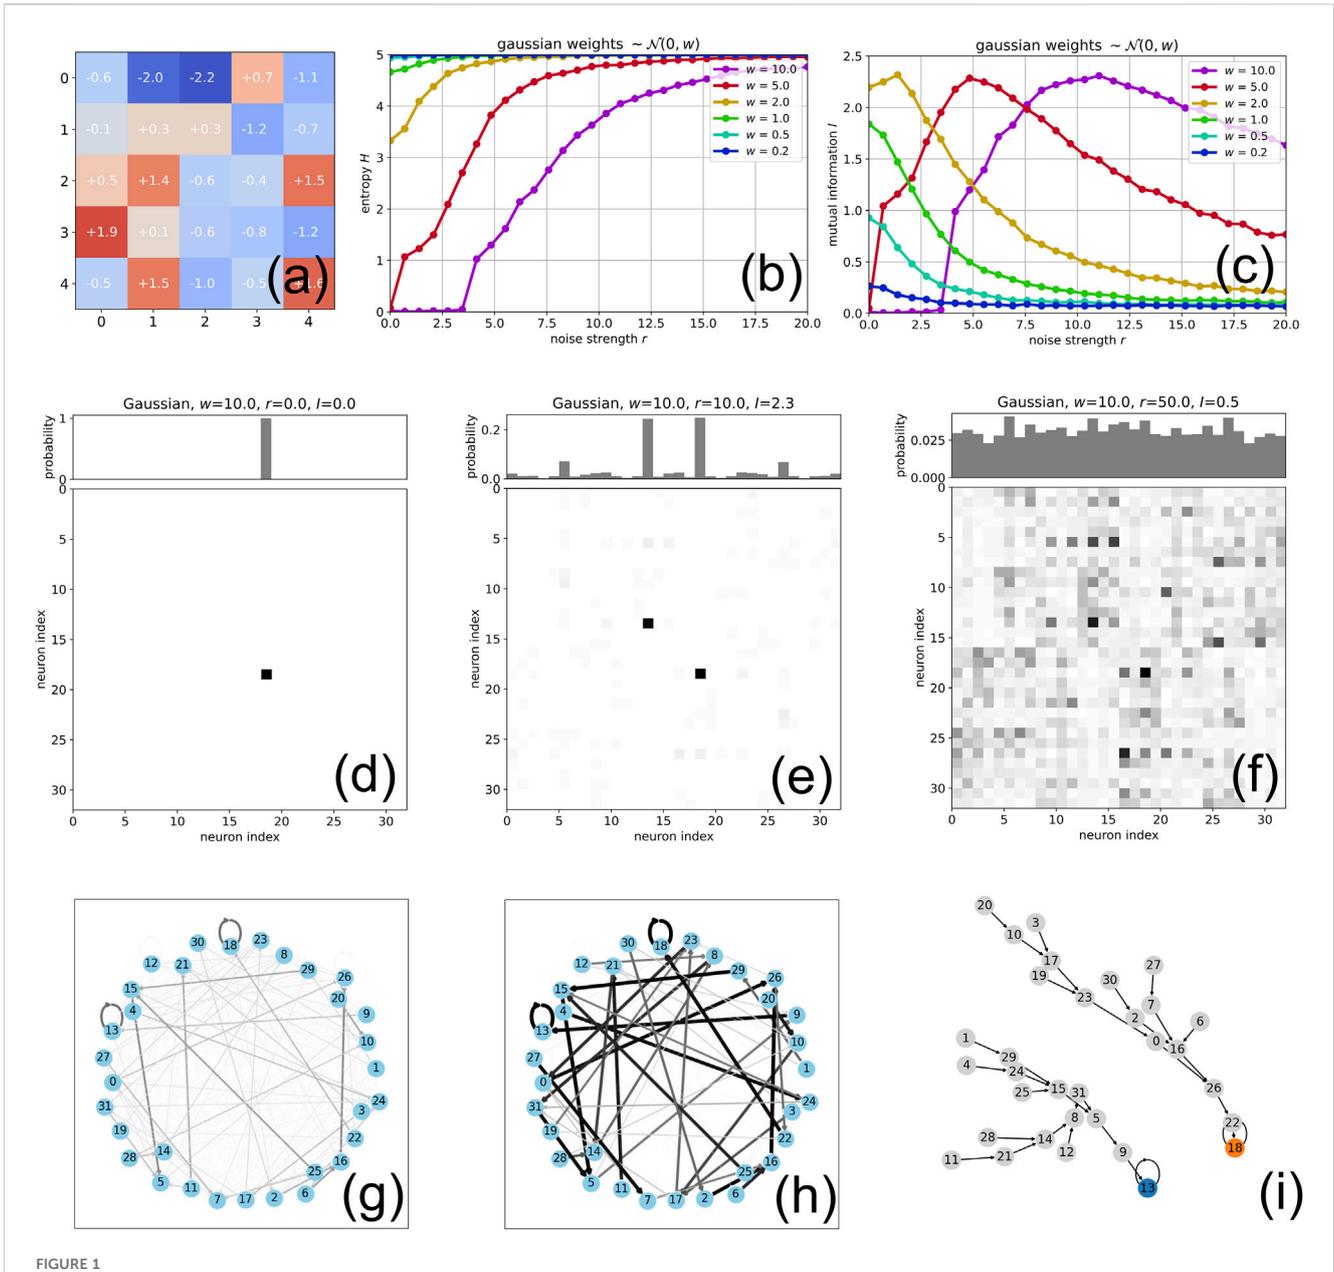


FIGURE 1 Effect of weight magnitude on Recurrence Resonance in Boltzmann Machines: **(A)** The elements of a 5×5 weight matrix are drawn randomly from a standard normal distribution $\sim \mathcal{N}(0, 1)$. The matrix is then scaled (multiplied) by a tunable weight magnitude parameter w , thus changing the standard deviation of the elements while keeping the fundamental structure of the matrix fixed. **(B)** The entropy H of global RNN states as a function of the strength r of added noise, for six different w . For very small weights $w \leq 1$ (dark blue, light blue and green curves) the system has near its maximum possible entropy $H = 5$, as the neurons fire almost independently. The entropy H at $r = 0$ is however drastically reduced as the weights get stronger ($w > 1$, orange, red and magenta curves), because the system is then trapped within a single attractor. In this trapped state, application of a sufficient level of noise $r > 0$ allows the system to escape and visit other attractors as well, leading to an increase of the entropy. **(C)** The mutual information I between subsequent global RNN states as a function of the strength r of added noise, for six different w . In the regime of low weights w , adding noise leads to a decrease of the mutual information I . However, for sufficiently large w (red and magenta curves), the mutual information I increases with added noise r and shows a maximum at some optimal noise level (at around $r_{opt} \approx 5$ for the red curve). **(D)** The joint probability $P(\mathbf{s}^{(t)}, \mathbf{s}^{(t+1)})$ of subsequent global system states, for the RNN with width magnitude parameter $w = 10$, without noise. The system is trapped in one of the 32 possible states, which is also visible in the marginal distribution $p(\mathbf{s}^{(t)})$ of system states (top of the matrix plot). This resting in the fixpoint corresponds to a mutual information of $I = 0$. **(E)** Adding a noise of strength $r = 10$ allows the system to visit another quasi-stable fixpoint (as well as some other transient states) for a fraction of time steps, increasing the mutual information to $I = 2.3$. **(F)** Adding a too large level $r = 50$ of noise opens up almost all possible states and transitions for the system. However, the system can no longer stay within any particular attractor for a longer period of time. This loss of order corresponds to a drop of the mutual information to $I = 0.5$. **(G, H)** State transition graph of the RNN for $w = 1$ **(G)** and $w = 5$ **(H)**, with transition probabilities coded by arrow thickness. As the weight magnitude w increases, preferred paths emerge, and thus an attractor landscape. **(I)** Graph of only the dominant state transitions at $w = 5$. The system has two fix points with large basins of attraction. For definition of state indices compare (Equation 6).

and the weighted directional edges represent the possible transitions $x \rightarrow y$ between pairs of states (Figures 1G, H). The probability of each transition, which is given by the conditional probability $P(y|x)$, may be indicated by the thickness of the edges. This state transition graph, which is determined only by the weight matrix \mathbf{W} , describes the *possible* dynamical behavior of the network completely, irrespective of which path the system is actually taking through the graph.

Below, we will visualize the aggregated activity in the network by the joint probability $P(y, x) = P(y|x)P(x)$. While this quantity also respects the fundamental transition possibilities (described by $P(y|x)$), it additionally accounts for which states and transitions the system has *actually* used during the observation time (described by $P(x)$).

3.1.1 Effect of increasing neural coupling without noise

We first consider the system without applying external noise ($r = 0$). For a weight magnitude of $w = 0$, the five neurons are completely isolated from each other and also have no self-connections (autapses). Consequently, each of the probabilistic SBM neurons produces a non-biased, temporally uncorrelated, binary random walk, where the two possible values -1 and $+1$ occur with equal probabilities, and independently from each other. This means that any momentary global system state x can transition to any other global successor state y with equal probability. We thus have a ‘structureless’, fully connected state transition graph, with equally ‘thick’ edges.

Since the time scale of $N_T = 10^4$ is long enough for the system to explore its entire space of $N_x = 32$ global states, we expect that for $w = 0$ the entropy reaches its maximum possible value $H = N_{neu} = 5$. As the system is purely random with a structureless state transition graph, the mutual information is expected to be $I = 0$. Consequently, the divergence $D = H - I$ is also maximal at $D = N_{neu} = 5$. This is indeed found in the numerical simulation.

As we tune the weight magnitude w from zero to increasingly positive values, the state transition graph, being determined by the weight matrix, is gradually developing a structure, that is, some of the transitions become more probable than others. Consequently, certain dominating paths are forming within the graph along the ‘thick’ edges, leading eventually to the emergence of dynamical attractors, such as fixed points (Figure 1I), higher period n -cycles, or transient states. These attractors are quite unstable at low weight magnitudes w , so that the system can still transition between them. Nevertheless, within the limited time horizon of N_T time steps, it now becomes impossible to visit all 32 global states with equal probability. Without additional noise ($r = 0$), this leads to a gradual decrease of the entropy $H(r = 0)$ (Figure 1B). The growth of structure in the graph with increasing w makes the system dynamics more deterministic and is thus also connected with a decrease of the random divergence $D(r = 5)$ (not shown in the figure).

However, since the entropy H and the divergence D decrease at different rates with the weight magnitude w , the mutual information $I = H - D$ shows a more complicated behavior (Figure 1C). As w is increased from 0.2 to 2, the mutual information without noise $I(r = 0)$ is first increasing, reflecting the higher degree of predictability of the next state. But for $w = 5$, we find that

$I(r = 0)$ is decreasing again, reaching a value of zero for $w = 10$ (Figure 1C, red and magenta curve).

This extreme situation of $H = D = I = 0$, which is most often found at large weight magnitudes but without noise, means that the system is trapped in a single global state, in other words, a quasi-stable fixed point x^* . Correspondingly, the joint probability $P(x, y)$ of subsequent system states has only a single non-zero entry at $x = y = x^*$ (Figure 1 (d, matrix plot)), and also the marginal state probability $P(x)$ has only a single entry at the fixed point state $x = x^*$ (Figure 1 (d, histogram on top)).

3.1.2 System behavior with noise

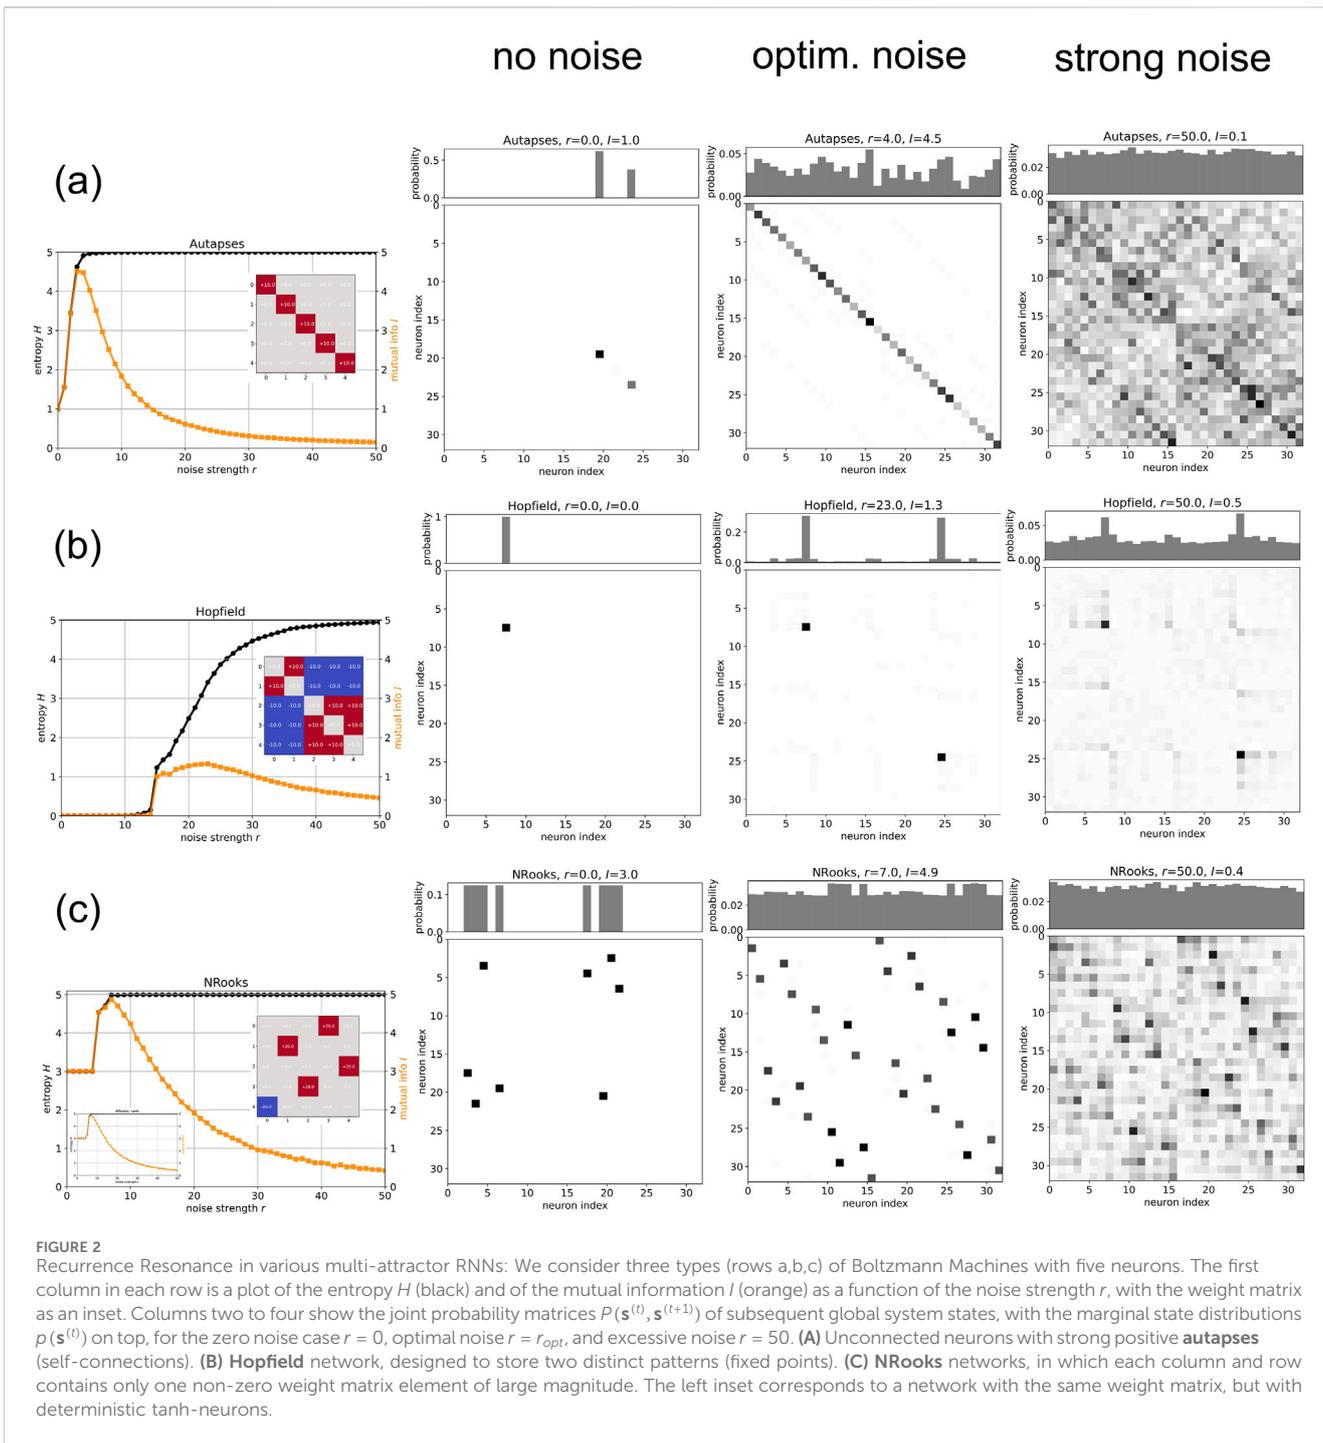
We now go back to the case of relatively weak weight magnitudes $w \leq 1$ and gradually increase the strength r of added noise. Here, the entropy H is already quite large without noise, and adding noise increases it even further (Figure 1 (b, green curves on the very top)). The noise thus helps the system to visit all possible states with equal probability, which can be seen as a beneficial effect. However, the noise increases also the divergence D at a fast rate (not shown in the figures), so that the mutual information $I = H - D$ is only decreasing as more noise is added (Figure 1 (c, green, light blue and dark blue curves)). We therefore do not observe RR, that is, a peak of the mutual information as a function of noise, in the regime of weak weight magnitudes.

In contrast, a different behavior is found for stronger weight magnitudes $w \geq 5$. Here, the entropy H is also increased by adding more noise (Figure 1 (b, red and magenta curves)), but the mutual information I is now initially increasing with noise - with a small exception at small noise levels (Figure 1 (c, red and magenta curves)). In the case of $w = 5$ (red curve), it reaches a maximum at a noise level of around $r = 5$ and then decreases again. For $w = 10$ (magenta curve), the maximum of I is around $r = 10$. Thus, RR is only observed in the regime of sufficiently strong weight magnitudes.

Generally, whenever the mutual information as a function of noise shows a clear maximum in a given network, the joint and marginal probability distributions are characteristically different at the points without noise (Figure 1D, $r = 0$), close to the RR maximum (Figure 1E, $r = 10$), and far beyond the RR maximum (Figure 1E, $r = 50$): Without noise, the system visits relatively few states, spending its time in only one or a few attractors. At the RR peak, the number of visited states is larger, and those states mainly belong to quasi-stable attractors (The matrix plot then shows a distinct set of dominating entries). Beyond the RR peak, the system visits even more states, but now only transiently, without staying in any particular attractor for a longer period of time (The matrix plot then appears more uniform and unstructured than at the RR maximum).

3.2 RR in selected networks with multiple attractors

If RR is a process where noise helps neural networks to reach more attractors in a given time horizon, the phenomenon should be particularly pronounced in systems with multiple (as well as sufficiently stable) attractors. We therefore select the following three specific types of weight matrices, while keeping the network



size of the SBM at $N_{neu} = 5$ and the time scale at $N_T = 10^4$: an autapses-only network, a Hopfield network, and a NRooks network.

3.2.1 Autapses-only network

We first test a ‘autapses-only’ network, in which all non-diagonal elements of the weight matrix (corresponding to inter-neuron connections) are zero, whereas the diagonal elements (corresponding to neuron self-connections, or ‘autapses’) have the same positive value $w = +10$ (Figure 2 (row (a), inset of left plot)). In this case, since the SBM neurons are isolated from each other, they produce mutually independent, binary random walks.

The probabilities of -1 and $+1$ are still equal, because we do not use biases ($v_n = 0$, See Methods). However, due to the excitatory autapses, the random walks are now temporally correlated (‘persistent’), which means that an output of $+1$ is more likely followed by another $+1$, and analogously for -1 . Each neuron thus tends to produce longer chains of outputs with identical sign, only switching to the opposite sign after a certain correlation time. For the total system, this means that any momentary global state is conserved, with high probability, for a finite number of time steps. Hence, each of the 32 global system states is a quasi-stable fixed point here, and the degree of stability can be increased by the self-

connection strength w . For this reason, this type of network could be useful as short-term memories in practical applications.

In our simulation, the autapse-only network without external noise is spending all of the 10^4 time steps in only two of its 32 fixed point attractors, consequently leading to a mutual information of about $I = 1$ (Figure 2 (row (a), column ‘no noise’)). Adding an optimal amount of noise ($r = 4$) drives the network to visit all available attractors, yet not with the same frequency. Nevertheless, the peak mutual information is with $I = 4.5$ close to the upper limit of 5 (Figure 2 (row (a), column ‘optim. noise’)). Applying an excessive noise of $r = 50$ lets the system undergo almost all 32^2 possible state-to-state transitions, but these frequent unpredictable jumps between attractors lead to a mutual information of only $I = 0.1$ (Figure 2 (row (a), column ‘strong noise’))

A remarkable feature of the autapse-only network’s RR-curve (Figure 2 (row (a))) is the part between zero and optimal noise. In this regime, the mutual information I follows extremely closely the rising curve of the entropy H , meaning that the divergence D is extremely small. Hence, the noise is on the one hand able to occasionally transfer the system from one (fixed point) attractor to a different one, but on the other hand allows the system to stay for a sufficiently long time interval within each attractor, so that the next state remains predictable to a high degree which constitutes a precondition for a high mutual information. We will see below that other systems also show this initial noise regime where entropy and mutual information rise together, while the divergence remains close to zero. At some level of noise, of course, the divergence must increase as well.

3.2.2 Hopfield network

Another type of recurrent neural network that is famous for its ability to have multiple (designable) fixed point attractors is the **Hopfield network** (Hopfield, 1982). Weight matrices of Hopfield networks are symmetric ($w_{mn} = w_{nm}$) and have no self-connections ($w_{mm} = 0$). Its neurons are traditionally updated one by one in an asynchronous manner, but we continue to use a synchronous update in our SBM model.

We have designed the weight matrix to ‘store’ the two patterns $[+1 + 1 - 1 - 1 - 1] \equiv 24$ and $[-1 - 1 + 1 + 1 + 1] \equiv 7$. The magnitude w of the weight matrix elements was made large enough to ensure a good stability of the two fixed points corresponding to the stored patterns (Figure 2 (row (b), inset of left plot)).

In a broad initial regime of noise strengths $r = 0 \dots 10$, the RR-curve of the Hopfield network (Figure 2 (row (b), left plot)) shows an entropy H and mutual information I of zero, which is characteristic for a system that is trapped in a single fixed point. The matrix plot of the joint probability reveals that this fixed point is the global state $[+1 + 1 - 1 - 1 - 1] \equiv 24$, the first of the two stored patterns (Figure 2 (row (b), column ‘no noise’)). The mutual information I starts to rise sharply at around $r = 15$ and reaches a peak at $r = 23$, while the entropy H continues to increase toward the upper limit. At the optimal noise level, the system is now visiting both fixed points (state 7 as well as state 24) with similar frequency, resulting in a peak mutual information of about $I = 1.3$ (Figure 2 (row (b), column ‘optim. noise’)). Since the two fixed points are very stable (due

to the large magnitude of matrix elements), the system is spending a comparatively large fraction of time in these two states, even at a large noise level of $r = 50$ (Figure 2 (row (b), column ‘strong noise’)).

3.2.3 NRooks network

Finally, we test the RR phenomenon in so-called ‘**NRooks networks**’, which under ideal conditions (large magnitude w of non-zero weights and long observation time scale N_T) are known to approach the upper theoretical limit of mutual information and entropy, corresponding to $I = H = N_{neu}$ (Metzner et al., 2024), and a vanishing divergence $D = 0$. The NRooks weight matrix has only one non-zero matrix element in each row and column (hence the name), and these N_{neu} non-zero matrix elements have the same magnitude w , but arbitrary signs (Figure 2 (row (c), upper inset of left plot)). It has been shown that all global states of an NRooks system are parts of n -cycles of various sizes, that is, there are no transient states that would merely lead into these attractors (Metzner et al., 2024).

Our specific NRooks system turns out to have four different 8-cycles as attractors, and without noise it is trapped in one of them (Figure 2 (row (c), column ‘no noise’)). Since running for thousands of time steps within this attractor involves eight distinct states (creating entropy) in a perfectly predictable order (without divergence D), the entropy and mutual information have already a relatively large value of $H = I = 3$, even without external noise (Figure 2 (row (c), left plot)). Since attractors are quite stable at $w = 20$, we observe a plateau with $H = I = 3$ in the RR-curve, holding up to a noise level of $r = 4$, where H and I start to increase rapidly. At the peak of the mutual information, occurring for a noise level of $r = 7$, it reaches the value of $I = 4.9$, which is close to the theoretical maximum of 5. Indeed, at that point the system is visiting all four 8-cycle attractors with about the same frequency (Figure 2 (row (c), column ‘optim. noise’)). It stays for a very long time in each of them, behaving almost perfectly deterministic. Only occasionally, the noise of optimal strength ‘kicks’ the system randomly to one of the other three attractors. As usual, for a very strong amount of noise, the system loses its predictability, and the mutual information I drops accordingly, while the entropy H remains at the upper limit.

In the above numerical experiments with multi-attractor SBMs, we have used relatively large weight magnitudes w . In the given context, this served the purpose to make the attractors of the autonomous networks more stable. More generally, large weight magnitudes drive the SBM neurons into the saturation regime of the logistic activation function, so that the on-probabilities $p_{on,n}^{(t)}$ become either ≈ 0 or ≈ 1 for all neurons n and all time steps t . Hence, the probabilistic SBM then behaves quasi deterministic. For this reason, we could use an SBM to implement a Hopfield network, which is usually based on deterministic binary threshold neurons.

In order to further demonstrate the saturation regime of the SBM, we have used the same weight matrix that was used in the NRooks example also in a network with deterministic tanh-Neurons (See Methods for details). The resulting RR-curve is indeed extremely similar to that of the probabilistic SBM (Figure 2 (row (c), lower inset of left plot)).

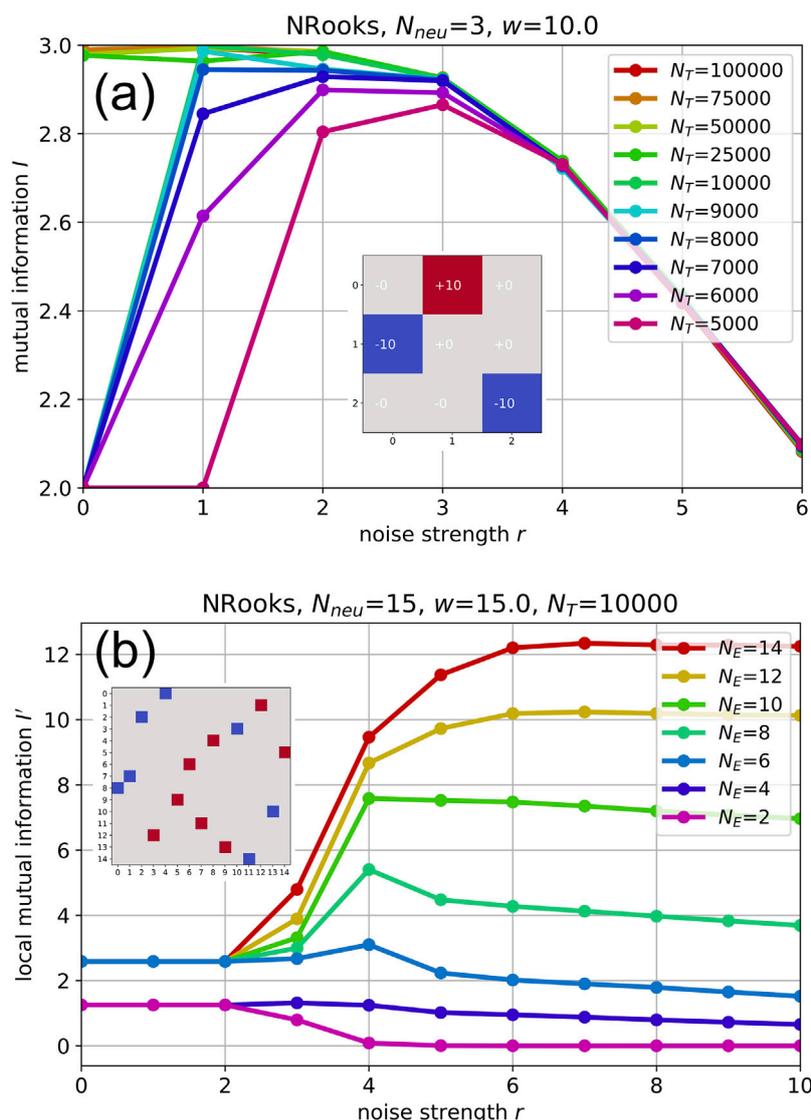


FIGURE 3
(A): Time-scale dependence of information quantities. Mutual information I in a 3-neuron NRooks system, with weight magnitude $w = 10$ of the non-zero matrix elements, as a function of the noise strength r , evaluated for different numbers of simulation time steps N_T between 5000 and 100,000. All curves show the same decay of mutual information I for noise levels larger than about $r = 4$, but they differ drastically in the regime of smaller noise levels: For a proper time scale $N_T = 5000$, the curve $I = I(r)$ shows clearly the RR phenomenon, with a peak mutual information at $r = 3$. As the observation time scale N_T is increased, the mutual information I in the whole regime of smaller noise levels is rising, as the system gets more opportunities to switch attractors. This leads to a shift of the RR peak to lower noised levels r_{opt} . For time scale above $N_T \approx 10000$, the mutual information $I(r = 0)$ without noise is rather abruptly jumping to close to the maximum possible value, meaning that the system does not require external help any more to realize optimal information flux. At this point the RR phenomenon disappears. **(B): 'Local' mutual information I'** in a 15-neuron NRooks system with weight magnitude $w = 15$ and time scale $N_T = 10^4$, as a function of the noise strength r , evaluated only in a subgroup of $N_E \in \{2, 4, 6, 8, 10, 12, 14\}$ neurons. Since the time scale of $N_T = 10^4$ is insufficient to reach 'ergodic' behaviour in a space of 2^{15} states, no RR maximum is visible for large subgroup sizes $N_E = 12$ and $N_E = 10$, but instead a rise and saturation of I' with increasing noise (plateau remains until at least $r = 50$, data not shown). For too small subgroup sizes $N_E = 2$ and $N_E = 4$, only a monotonous drop of mutual information is observed as a function of the noise level r . However, for intermediate subgroup size $N_E = 6$ and $N_E = 8$, a clear RR peak is emerging. It is therefore possible to detect RR even in large systems, at least locally.

3.3 Time-scale dependence of RR

As already mentioned above, the observation time scale N_T is another critical factor that determines whether a RR peak will be observable in a given system.

For a demonstration, we now use a NRooks system of only three neurons (Figure 3A). Because of its extremely small state space ($N_x = 8$), it is possible to make the attractors stable by a relatively

large weight magnitude of $w = 10$, but nevertheless to approach the ergodic limit of very large time scales N_T , where the system is able to visit all its attractors autonomously, without the injection of external noise.

We find that for too strong levels of noise (here $r \geq 4$), the mutual information I is just declining, irrespective from the time scale N_T . In this regime, the system is operating already at maximum entropy ($H \approx 3$), but the noise is causing an increasing loss of predictability.

For a time scale of $N_T = 5000$, which is appropriate for the given system, a clear RR peak is observed in the mutual information curve $I = I(r)$ at around $r = 3$. As the time scale is now prolonged, up to around $N_T = 9000$, the mutual information $I = I(r)$ is generally rising for moderate noise levels in the range $r \in [0, 3]$, because the systems gets more opportunities to escape and switch from one attractor state to another. As a consequence, the RR peak is moving to smaller noise levels r_{opt} .

For the moderate time scales $N_T < 9000$ considered so far, the system, without noise, is trapped in a 4-cycle, and therefore the mutual information is $I(r = 0) = 2$. However for larger time scales $N_T > 10000$, we find a value $I(r = 0) \approx 3$ close to the theoretical maximum of 3, which means that now even the zero-noise system can visit all its attractor states and run through each of them in an almost perfectly predictable way. Thus, the phenomenon of RR is not observable on extremely long time scales, where systems already operate close to their ergodic regime.

3.4 “Local” mutual information in sub-networks

In actual applications of RNNs, such as reservoir computing, networks are typically so large ($N_{neu} > 100$) and consequently the state spaces so huge ($N_x > 2^{100}$) that the ergodic regime cannot be reached on any practical time scale N_T . Moreover, in such practically non-ergodic systems, it is also impossible to accurately evaluate the full mutual information of subsequent system states, because the joint probability matrices are too large ($N_x \times N_x > 2^{100} \times 2^{100}$) and because the empirical distributions have not enough time to converge toward a stable result. The question then arises how to compute a useful approximation of $I = I(r)$ in large systems, even when they are observed on ‘too short’ (but practically relevant) time scales.

Although a detailed investigation of this question is beyond the scope of the present paper, we provide a first insight using a 15-neuron NRooks system, observed on the non-ergodic time scale of $N_T = 10000$ (Figure 3B). To alleviate the matrix size problem, we only consider global states that have actually been visited during the observation time (See Methods for details). Moreover, as a proxy for the full $I(r)$, we compute the ‘local’ mutual information $I'(r)$ within smaller sub-networks, i.e., subgroups of only $N_E \leq N_{neu}$ neurons.

For large sub-networks ($N_E > 10$), instead of a RR peak, we find an increase of the local mutual information $I'(r)$ with noise, and finally a saturation. This plateau is also observed for larger noise levels up to $r = 50$ (data not shown).

In contrast, for very small sub-networks ($N_E \leq 4$), we find already a smaller starting value of $I'(r = 0)$ at zero noise, and eventually a decline of $I'(r)$ with increasing noise level r .

However, for a certain intermediate range of sub-network sizes between $N_E = 6$ and $N_E = 8$, the curve $I'(r)$ shows a clear maximum that decays very slowly after the peak. Thus, in large networks, when observed on short quasi non-ergodic time scales, a phenomenon similar to RR can occur within smaller sub-networks, whereas the mutual information of the total system then shows a saturation-type dependence on the noise level.

3.5 Effect of noise pulses in probabilistic, binary-valued RNNs

At the peak of the RR curve, the continuous white noise input of optimal strength r_{opt} is leaving a RNN in its present attractor for long times, but occasionally causes a random transit to one of the other available attractors. It is this combination of high predictability and high entropy that leads to the optimal value of the mutual information.

A natural extension (and putative application) of this concept are short noise pulses - applied only at times when a change of attractor state is required - instead of a continuous feed-in of noise. To test this concept, we have again used the 5-neuron NRooks system of Figure 2C, with its four different 8-cycles as attractors. The system is initially in one of its 8-cycle attractors (Figure 4A), and remains in this attractor for an arbitrarily long period that only depends on the weight magnitude w . By applying short (10 time steps) yet strong ($r = 50$) Gaussian white noise pulses, we could indeed transfer the system randomly to one of the other attractors. It also happens that the system ends up in the same attractor (yet at a different ‘phase’ of the periodic cycle), but eventually we could reach all four 8-cycles by this way (Figures 4B–D).

In actual applications, the noise strength r could be made a continuously or abruptly changing function of time, designed to accelerate the ‘equilibration’ of the RNN over its state space. This clearly resembles the well-known technique of simulated annealing (Kirkpatrick et al., 1983; Aarts et al., 1987; Aarts and Van Laarhoven, 1989), yet with the important difference that quasi-deterministic RNNs follow a prescribed order of states, once they have entered a cyclic attractor. In biological neural networks, the optimal momentary noise strength might even be provided by a feedback control system that continuously aims to optimize network performance.

3.6 Effect of noise pulses in continuous-valued RNNs

So far, we have only briefly explored the effect of noise on networks of deterministic tanh-neurons (inset of Figure 2C). As a further glimpse into this alternative field of research, we apply a short (5 time steps) and weak ($r = 5$) noise pulse to a very small (3 neurons) tanh-network, in which the nine matrix elements have been drawn randomly from a standard normal distribution $\sim \mathcal{N}(0, 1)$.

Before the noise pulse, the system is allowed to run freely for 100 time steps. The resulting system states at each time step are here continuous points within the three-dimensional cube $]-1, +1[^3$ and can thus be visualized directly as a 3-dimensional trajectory (Figures 4E–G). We find the system initially within a ‘strange’, loop-like attractor (e). During the short noise pulse, the trajectory is erratic and reaches the borders of the state space cube (f). After the pulse, the system has settled in a new strange attractor, which resembles a 2-cycle, but only with an approximate return to the end points in each oscillation period. Thus, it is also possible to achieve a switch of attractor states in deterministic RNNs with continuous output values by the injection of noise pulses.

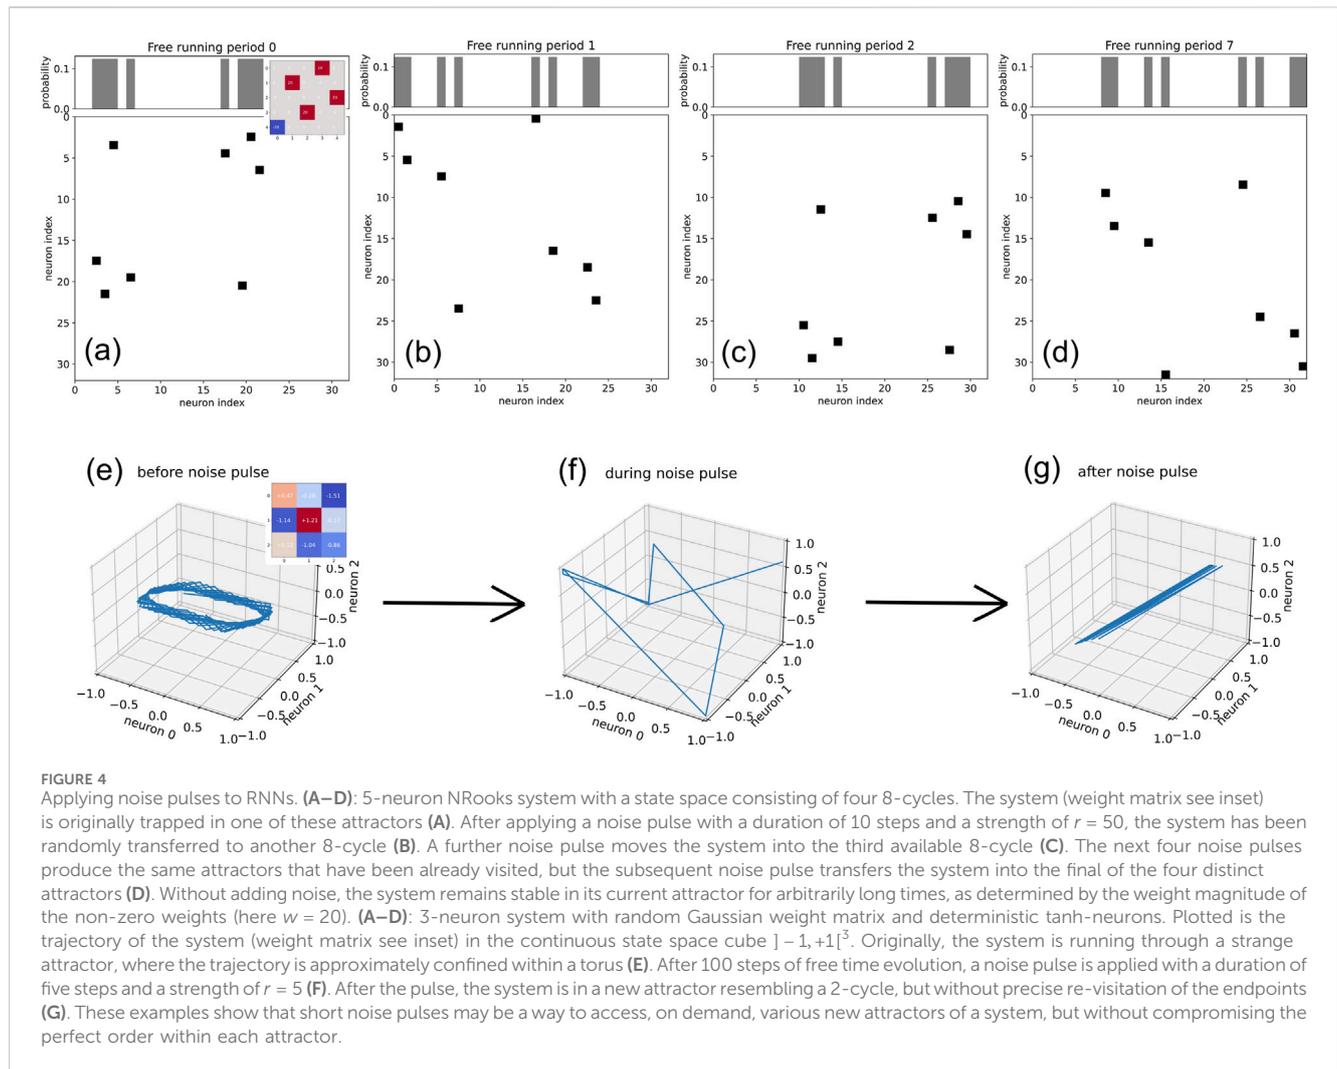


FIGURE 4

Applying noise pulses to RNNs. (A–D): 5-neuron NRooks system with a state space consisting of four 8-cycles. The system (weight matrix see inset) is originally trapped in one of these attractors (A). After applying a noise pulse with a duration of 10 steps and a strength of $r = 50$, the system has been randomly transferred to another 8-cycle (B). A further noise pulse moves the system into the third available 8-cycle (C). The next four noise pulses produce the same attractors that have been already visited, but the subsequent noise pulse transfers the system into the final of the four distinct attractors (D). Without adding noise, the system remains stable in its current attractor for arbitrarily long times, as determined by the weight magnitude of the non-zero weights (here $w = 20$). (A–D): 3-neuron system with random Gaussian weight matrix and deterministic tanh-neurons. Plotted is the trajectory of the system (weight matrix see inset) in the continuous state space cube $[-1, +1]^3$. Originally, the system is running through a strange attractor, where the trajectory is approximately confined within a torus (E). After 100 steps of free time evolution, a noise pulse is applied with a duration of five steps and a strength of $r = 5$ (F). After the pulse, the system is in a new attractor resembling a 2-cycle, but without precise re-visitation of the endpoints (G). These examples show that short noise pulses may be a way to access, on demand, various new attractors of a system, but without compromising the perfect order within each attractor.

4 Conclusion

In this work, we have re-considered the phenomenon of Recurrence Resonance (RR), i.e., the peak-like dependence of an RNN's internal information flux on the level r of white noise added to each of the neurons (Krauss et al., 2019a). The information flux is measured by the mutual information $I = H - D$ between subsequent system states, a quantity that grows as more states become available (larger entropy H), and/or when each successor state can be better predicted from its predecessor (smaller divergence D).

We have shown that a resonance-like peak of $I(r)$ can only be observed in networks that fundamentally have a whole set of relatively stable dynamical attractors available, but which - without external intervention - would remain trapped in one of them during the entire observation time scale N_T . In this situation, adding a small level $r \leq r_{opt}$ of noise helps such networks to occasionally jump out of the present attractor and switch into another one, without significantly reducing the predictability of the state sequence within each of these quasi-stable attractors (strong increase of H , but weak increase of D). If the noise level r is however increased beyond the optimal point r_{opt} , predictability is lost and consequently the information flux $I(r)$ is declining again,

while the entropy H is still increasing toward its upper limit. By contrast, networks that already have a high internal information flux from the beginning will not show a RR peak, but only a decline of I as a function of r .

We have demonstrated the RR phenomenon using Symmetric Boltzmann Machines (SBMs) with different types of weight matrices, including random Gaussian matrices, diagonal matrices (autapse-only networks), Hopfield networks trained on specific patterns, and in NRooks systems that are known to reach the upper limit of information flux. In each case, we demonstrated that the network without noise is trapped in a single or few attractors, based on the joint probability of subsequent system states. An optimal level of noise makes more (or even all) attractors available without too many unpredictable transitions. However, excessive levels of noise cause more or less random jumps between all possible pairs of states. In systems with a very high stability of attractors (induced by a large weight magnitude w), we have found that $I(r)$ remains constant at the initial value $I(r = 0)$ for a certain range of noise levels, before it abruptly rises in the way of a phase transition.

We have also demonstrated that RR can only be observed in appropriate time scales N_T , relative to the total number of possible

system states $N_x = 2^{N_{neu}}$. For arbitrarily long observation times (which are of theoretical interest but not so much of practical relevance), a neural network operating in the quasi-deterministic regime, but with at least a small probabilistic component (like SBMs with a large but finite weight magnitude w) can eventually visit all its attractors in an ergodic manner, but still stay in each of them for extended time intervals. Then $I(r=0)$ is already close to the optimum value and application of noise only degrades the information flux.

An interesting problem arises therefore in networks with many neurons and thus an exponentially large state space, such as reservoir computers, or brains. Such systems will necessarily spend all their lives within a negligible fraction of the fundamentally available state space, possibly consisting of only a tiny subset of attractors. One way to cope with this ‘practical non-ergodicity’ would be a repeated active switching between attractor subsets (perhaps using noise pulses), until a useful one is found, and then to stay there. Alternatively, the networks may be designed (or optimized) such that the useful attractors have a very large basin of attraction. A similar problem has been discussed in the context of protein folding with the ‘Levinthal paradox’, where naturally existing proteins fold into the desired conformation much faster than expected by a random thermal search in conformation space (Zwanzig et al., 1992; Karplus, 1997; Honig, 1999), probably due to funnel-type energy landscapes (Bryngelson et al., 1995; Martínez, 2014; Wolynes, 2015; Röder et al., 2019).

In our context of the RR phenomenon, practical non-ergodicity makes it impossible to compute the stationary information flux in a large network, because the system never reaches a stationary state (marked by constant probability distributions) within any practical time scale N_T . When the information flux in a network is evaluated naively, using the ‘transient’ (not yet converged) joint probability distributions, we have found that $I(r)$ shows a saturating behavior, rather than a maximum. Nevertheless, the ‘local’ mutual information, evaluated for a suitably sized sub-network, can then still show a RR peak.

Finally, we have explored the repeated application of short noise pulses, rather than feeding continuous noise into the neurons. We could demonstrate that each pulse offers the network a chance to switch to a new random attractor (such as an n-cycle), while the intermediate free running phases allow the system to deterministically and thus predictably follow the fixed order of states within each of the attractors.

If the random noise signals are delivered independently to each individual neuron, and if the levels of these external control signals are strong enough to override the recurrent internal signals from the other neurons, then the network could theoretically end up, after the noise application, in any of its 2^n global system states. Thus all the system’s attractors could theoretically be accessed by this way. Nevertheless, a complete exploration of the entire attractor landscape will take an extremely long time in large networks.

Based on this noise-induced random switching mechanism, an evolutionary optimization algorithm could be implemented in a recurrent neural network, in which various attractors are tried until one turns out useful for a given task. We speculate that this principle might be used in central pattern generators of biological brains (Hooper, 2000), for example, in order to find temporal activation patterns for certain motor tasks (Marder and Bucher, 2001).

5 Discussion

In free running SBMs, a neuron’s probability of being ‘on’ in the next time step is computed by a logistic activation function $\sigma(u)$, where u is the weighted sum of inputs from other neurons. Adding white, normally distributed noise to u corresponds to a convolution of the logistic function with a Gaussian kernel, resulting in a ‘broadening’ of the activation function. Increasing the noise level thus has an effect similar to turning up the ‘temperature’ parameter T in a re-scaled activation function $\sigma(u/T)$. This opens up a new interpretation of the RR phenomenon in terms of statistical physics, in particular if an energy $E(\mathbf{s}) = -\sum_{mn} w_{mn} s_m s_n$ can be assigned to each global system state \mathbf{s} .

In this case, the system - after sufficiently long time - would come to thermal equilibrium, and the probability of finding it at any state \mathbf{s} would be proportional to the Boltzmann distribution $\propto e^{-E(\mathbf{s})/T}$. At low temperature, the system would therefore spend most of its time in the deepest valleys of the energy landscape, which may correspond to fixed point attractors. Increasing the temperature would lead to a more uniform distribution of states over the energy landscape, and there might be a sweet spot for the temperature that perfectly balances stability of the attractors with occasional barrier crossings, just as in the RR phenomenon. Moreover, during the equilibration process, it might be advantageous to start with a high temperature (noise level) and then to slowly decrease it, as in Simulated Annealing (Van Laarhoven et al., 1987; Bertsimas and Tsitsiklis, 1993).

Note, however, that in our SBM model the system’s dynamics cannot be visualized as a simple probabilistic downhill relaxation within the energy landscape $E(\mathbf{s})$, because we are using a synchronous update of all neurons and non-symmetric weights. For example, in NRooks systems, each n-cycle attractor can have another energy, but the energy is the same for all states that belong to the same attractor. Despite of this ‘energetical degeneration’ of states within a given n-cycle, the system is not randomly jumping between those states, as it would be expected from a thermal system, but it is running through the state sequence in a perfectly deterministic way.

In this work, we have mainly focused on probabilistic SBMs as model systems of recurrent neural networks. However, we have shown that for sufficiently large weight magnitudes w the neurons operate in the saturation regime of the logistic activation function and therefore the SBMs behave quasi-deterministic. In this regime, the RR curves of SBMs turned out to be extremely similar to those of networks with the same weight matrix, but with deterministic tanh-neurons. Also, we have demonstrated that attractor switching by noise pulses works equally well for tanh-networks with continuous outputs. Nevertheless, regarding the details of the RR phenomenon, we expect future work to reveal some fundamental differences between probabilistic and deterministic RNNs. In particular, while the SBM neurons turn into independent random generators when the general weight magnitude w of their connections is turned down, deterministic networks can - even for small w - produce complex dynamical attractor states. It is not clear at present how sensitive those attractors react on externally injected noise.

Neural networks, both artificial and biological, have a tendency to become trapped in low-entropy dynamical attractors, which correspond to repetitive, predictable patterns of activity (Khona

and Fiete, 2022). These attractors are often associated with stable cognitive states or established perceptual interpretations (Beer and Barak, 2024). In particular, it has been shown that during spontaneous activity the brain does not randomly change between all theoretically possible states, but rather samples from the realm of possible sensory responses (Luczak et al., 2009; Schilling et al., 2024). This indicates that the brain's spontaneous activity encompasses a spectrum of potential responses to stimuli, effectively preconfiguring the neural landscape for incoming sensory information. While stability and predictability are crucial for efficient functioning and reliable behavior, they can also limit the flexibility and adaptability of the system (Khona and Fiete, 2022; Beer and Barak, 2024). This is particularly problematic in contexts requiring learning, creativity, and the generation of novel ideas (Sandamirskaya, 2013).

The introduction of noise into neural networks has been suggested as a mechanism to overcome the limitations imposed by these low-entropy attractors (Hinton and Van Camp, 1993). Noise, in this context, refers to stochastic fluctuations that perturb the network's activity, pushing it out of stable attractor states and into new regions of the state space (Bishop, 1995). This process can enhance the network's ability to explore a wider range of potential states (Sietsma and Dow, 1991), thereby increasing its entropy and promoting the discovery of novel solutions or interpretations.

Biological neural systems, such as the human brain, provide compelling evidence for the utility of noise in cognitive processes. The brain is inherently noisy, with intrinsic fluctuations occurring at multiple levels, from ion channel gating to synaptic transmission and neural firing (Faisal et al., 2008). This noise is not merely a byproduct of biological imperfection; rather, it plays a functional role in various cognitive tasks (McDonnell and Ward, 2011). For example, noise-induced variability in neural firing can enhance sensory perception by enabling the brain to detect weak signals that would otherwise be drowned out by deterministic activity (Deco et al., 2009).

Another noise-based phenomenon of great importance in physical and biological systems is Stochastic Resonance (Gammaitoni et al., 1998; Moss et al., 1993). It is typically occurring in signal detection systems, where the incoming signal needs to exceed a minimal threshold in amplitude to be detected. Adding an appropriate level of noise to the input can then stochastically lift even weak signal above the threshold and thereby improve the detection performance, as it has indeed been observed in various sensory systems (Moss et al., 2004; Stein et al., 2005; Ward, 2013). The phenomenon of Recurrence Resonance discussed in this paper is different from Stochastic Resonance, as it improves the spontaneous, internal information flux in a neural network, rather than the signal transmission from the outside to the inside of the system. However, using mutual information or correlation-based measures, it is also possible to quantify the information flux from the input nodes of a RNN at time t to the internal network state at a later time $t+\Delta t$ (Metzner and Krauss, 2022).

Noise also facilitates learning and plasticity. During development, random fluctuations in neural activity contribute to the refinement of neural circuits, allowing for the fine-tuning of synaptic connections based on experience (Marzola et al., 2023; Zhang et al., 2021; Fang et al., 2020). In adulthood, noise can help the

brain escape from local minima during learning processes, thereby preventing overfitting to specific patterns and promoting generalization (Zhang et al., 2021; Fang et al., 2020). This is particularly relevant in the context of reinforcement learning, where exploration of the state space is crucial for finding optimal strategies (Weng, 2020; Bai et al., 2023).

Moreover, noise-induced transitions between attractor states can support cognitive flexibility and creativity. For instance, the ability to switch between different interpretations of ambiguous stimuli (Panagiotaropoulos et al., 2013), or to generate novel ideas, relies on the brain's capacity to break free from dominant attractor states and explore alternative possibilities (Wu and Koutstaal, 2020; Jaimes-Reátegui et al., 2022). This is consistent with the observation that certain cognitive disorders, characterized by rigidity and a lack of flexibility (e.g., autism, obsessive-compulsive disorder), are associated with reduced neural noise and hyper-stable attractor dynamics (Dwyer et al., 2024; Watanabe et al., 2019).

In conclusion, the investigations detailed in our study firmly establish Recurrence Resonance (RR) as a genuine emergent phenomenon within neural dynamics: the mutual information of the system is increased by the addition of noise that itself has zero mutual information. Hence, the application of optimal noise levels can transform neural systems from states of minimal information processing capabilities to significantly enhanced states where information flow is not only possible but also maximized. This effect, whereby noise beneficially modifies system dynamics, underscores the complex and non-intuitive nature of neural information processing, presenting noise not merely as a disruptor but as a critical facilitator of dynamic neural activity. This finding opens up new avenues for exploiting noise in the design and enhancement of neural network models, particularly in areas demanding robust and adaptive information processing.

The introduction of noise into neural networks can be seen as a fundamental mechanism by which the brain enhances its cognitive capabilities. By destabilizing low-entropy attractors and promoting the exploration of new states, noise enables learning, perception, and creativity. This perspective not only aligns with empirical findings from neuroscience but also offers a theoretical framework for understanding how complex cognitive functions can emerge from the interplay between deterministic and stochastic processes in neural systems.

Furthermore, the insights gained from our study provide a valuable foundation for advancing artificial intelligence (AI) technologies, particularly in the realms of reservoir computing and machine learning. Reservoir computing, which leverages the dynamic behavior of recurrent neural networks, can benefit from the strategic introduction of noise to enhance its computational power and adaptability. Similarly, machine learning models can incorporate noise to avoid overfitting, explore diverse solution spaces, and improve generalization. By integrating these principles, AI systems can emulate the brain's ability to learn and adapt in complex, unpredictable environments, leading to more robust and innovative technological solutions. This convergence of neuroscience and AI not only deepens our understanding of cognitive processes but may also drive the development of next-generation intelligent systems capable of solving real-world problems with unprecedented efficiency and creativity.

Data availability statement

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

Author contributions

CM: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Software, Validation, Visualization, Writing—original draft, Writing—review and editing. AS: Funding acquisition, Validation, Writing—original draft. AM: Resources, Supervision, Validation, Writing—original draft. PK: Conceptualization, Funding acquisition, Methodology, Project administration, Resources, Supervision, Validation, Writing—original draft, Writing—review and editing.

Funding

The author(s) declare that financial support was received for the research, authorship, and/or publication of this article. This work was funded by the Deutsche Forschungsgemeinschaft (DFG,

German Research Foundation): KR 5148/3-1 (project number 510395418), KR 5148/5-1 (project number 542747151), and GRK 2839 (project number 468527017) to PK, and grant SCHI 1482/3-1 (project number 451810794) to AS.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

The author(s) declared that they were an editorial board member of Frontiers, at the time of submission. This had no impact on the peer review process and the final decision.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

References

- Aarts, E. H. (1987). Simulated annealing: theory and applications. *Reidel*.
- Aarts, E. H. L., and Van Laarhoven, P. J. M. (1989). Simulated annealing: an introduction. *Stat. Neerl.* 43 (1), 31–52. doi:10.1111/j.1467-9574.1989.tb01245.x
- Alzubaidi, L., Zhang, J., Humaidi, A. J., Al-Dujaili, A., Duan, Ye, Al-Shamma, O., et al. (2021). Review of deep learning: concepts, cnn architectures, challenges, applications, future directions. *J. big Data* 8 (1), 53–74. doi:10.1186/s40537-021-00444-8
- Bai, C., Liu, P., Liu, K., Wang, L., Zhao, Y., Han, L., et al. (2023). Exploration in deep reinforcement learning: from single-agent to multiagent domain. *IEEE Trans. Neural Netw. Learn. Syst.* 34, 4776–4790. doi:10.1109/tnnls.2021.3129160
- Barak, O. (2017). Recurrent neural networks as versatile tools of neuroscience research. *Curr. Opin. Neurobiol.* 46, 1–6. doi:10.1016/j.conb.2017.06.003
- Beer, C., and Barak, O. (2024). Revealing and reshaping attractor dynamics in large networks of cortical neurons. *PLOS Comput. Biol.* 20 (1), e1011784. doi:10.1371/journal.pcbi.1011784
- Bertschinger, N., and Natschläger, T. (2004). Real-time computation at the edge of chaos in recurrent neural networks. *Neural Comput.* 16 (7), 1413–1436. doi:10.1162/089976604323057443
- Bertsimas, D., and Tsitsiklis, J. (1993). Simulated annealing. *Stat. Sci.* 8 (1), 10–15. doi:10.1214/ss/1177011077
- Bishop, C. M. (1995). Training with noise is equivalent to tikhonov regularization. *Neural Comput.* 7 (1), 108–116. doi:10.1162/neco.1995.7.1.108
- Boedecker, J., Obst, O., Lizier, J. T., Mayer, N. M., and Asada, M. (2012). Information processing in echo state networks at the edge of chaos. *Theory Biosci.* 131 (3), 205–213. doi:10.1007/s12064-011-0146-8
- Bönsel, F., Krauss, P., Metzner, C., and Yamakou, M. E. (2022). Control of noise-induced coherent oscillations in three-neuron motifs. *Cogn. Neurodynamics* 16 (4), 941–960.
- Brunel, N. (2016). Is cortical connectivity optimized for storing information? *Nat. Neurosci.* 19 (5), 749–755. doi:10.1038/nn.4286
- Bryngelson, J. D., Onuchic, J. N., Socci, N. D., and Wolynes, P. G. (1995). Funnels, pathways, and the energy landscape of protein folding: a synthesis. *Proteins Struct. Funct. Bioinforma.* 21 (3), 167–195. doi:10.1002/prot.340210302
- Büsing, L., Schrauwen, B., and Legenstein, R. (2010). Connectivity, dynamics, and memory in reservoir computing with binary and analog neurons. *Neural Comput.* 22 (5), 1272–1311. doi:10.1162/neco.2009.01-09-947
- Dambre, J., Verstraeten, D., Schrauwen, B., and Massar, S. (2012). Information processing capacity of dynamical systems. *Sci. Rep.* 2 (1), 514–517. doi:10.1038/srep00514
- Deco, G., Rolls, E. T., and Romo, R. (2009). Stochastic dynamics as a principle of brain function. *Prog. Neurobiol.* 88 (1), 1–16. doi:10.1016/j.pneurobio.2009.01.006
- Dwyer, P., Vukusic, S., Williams, Z. J., Saron, C. D., and Rivera, S. M. (2024). “neural noise” in auditory responses in young autistic and neurotypical children. *J. autism Dev. Disord.* 54 (2), 642–661. doi:10.1007/s10803-022-05797-4
- Faisal, A. A., Selen, L. P. J., and Wolpert, D. M. (2008). Noise in the nervous system. *Nat. Rev. Neurosci.* 9 (4), 292–303. doi:10.1038/nrn2258
- Fang, Y., Yu, Z., and Chen, F. (2020). Noise helps optimization escape from saddle points in the synaptic plasticity. *Front. Neurosci.* 14, 343. doi:10.3389/fnins.2020.00343
- Farrell, M., Recanatani, S., Moore, T., Lajoie, G., and Shea-Brown, E. (2022). Gradient-based learning drives robust representations in recurrent neural networks by balancing compression and expansion. *Nat. Mach. Intell.* 4 (6), 564–573. doi:10.1038/s42256-022-00498-0
- Folli, V., Gosti, G., Leonetti, M., and Ruocco, G. (2018). Effect of dilution in asymmetric recurrent neural networks. *Neural Netw.* 104, 50–59. doi:10.1016/j.neunet.2018.04.003
- Gammaitoni, L., Hänggi, P., Jung, P., and Marchesoni, F. (1998). Stochastic resonance. *Rev. Mod. Phys.* 70 (1), 223–287. doi:10.1103/revmodphys.70.223
- Gerum, R. C., Erpenbeck, A., Krauss, P., and Schilling, A. (2020). Sparsity through evolutionary pruning prevents neuronal networks from overfitting. *Neural Netw.* 128, 305–312. doi:10.1016/j.neunet.2020.05.007
- Gonon, L., and Ortega, J.-P. (2021). Fading memory echo state networks are universal. *Neural Netw.* 138, 10–13. doi:10.1016/j.neunet.2021.01.025
- Haruna, T., and Nakajima, K. (2019). Optimal short-term memory before the edge of chaos in driven random recurrent networks. *Phys. Rev. E* 100 (6), 062312. doi:10.1103/physreve.100.062312
- Haviv, D., Rivkind, A., and Barak, O. (2019). “Understanding and controlling memory in recurrent neural networks,” in *International conference on machine learning* (Long Beach, CA, United States: PMLR), 2663–2671.
- Hinton, G. E., and Van Camp, D. (1993). “Keeping the neural networks simple by minimizing the description length of the weights,” in *Proceedings of the sixth annual conference on computational learning theory*, 5–13.
- Honig, B. (1999). Protein folding: from the Levinthal paradox to structure prediction. *J. Mol. Biol.* 293 (2), 283–293. doi:10.1006/jmbi.1999.3006
- Hooper, S. L. (2000). Central pattern generators. *Curr. Biol.* 10 (5), R176–R179. doi:10.1016/s0960-9822(00)00367-5
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci.* 79 (8), 2554–2558. doi:10.1073/pnas.79.8.2554
- Ichikawa, K., and Kaneko, K. (2021). Short-term memory by transient oscillatory dynamics in recurrent neural networks. *Phys. Rev. Res.* 3 (3), 033193. doi:10.1103/physrevresearch.3.033193
- Ikemoto, S., DallaLibera, F., and Hosoda, K. (2018). Noise-modulated neural networks as an application of stochastic resonance. *Neurocomputing* 277, 29–37. doi:10.1016/j.neucom.2016.12.111

- Jaeger, H. (2001). *The "echo state" approach to analysing and training recurrent neural networks-with an erratum note*. Bonn, Germany: German National Research Center for Information Technology GMD Technical Report, 148.
- Jaeger, H. (2014). Controlling recurrent neural networks by conceptors. *arXiv Prepr. arXiv:1403.3369*.
- Jaimes-Reategui, R., Huerta-Cuellar, G., García-López, J. H., and Pisarchik, A. N. (2022). Multistability and noise-induced transitions in the model of bidirectionally coupled neurons with electrical synaptic plasticity. *Eur. Phys. J. Special Top.* 231, 255–265. doi:10.1140/epjs/s11734-021-00349-w
- Kadmon, J., and Sompolinsky, H. (2015). Transition to chaos in random neuronal networks. *Phys. Rev. X* 5 (4), 041030. doi:10.1103/physrevx.5.041030
- Kaneko, K., and Suzuki, J. (1994). "Evolution to the edge of chaos in an imitation game," in *Artificial life III* (Citeseer).
- Karplus, M. (1997). The levinthal paradox: yesterday and today. *Fold. Des.* 2, S69–S75. doi:10.1016/s1359-0278(97)00067-9
- Katz, L., and Sobel, M. (1972). "Coverage of generalized chess boards by randomly placed rooks," in *Proceedings of the sixth berkeley symposium on mathematical statistics and probability, probability theory* (University of California Press), 6.3, 555–564. doi:10.1525/9780520375918-031Contributions to Probability Theory
- Khona, M., and Fiete, I. R. (2022). Attractor and integrator networks in the brain. *Nat. Rev. Neurosci.* 23 (12), 744–766. doi:10.1038/s41583-022-00642-0
- Kirkpatrick, S., Gelatt, C. D., Jr, and Vecchi, M. P. (1983). Optimization by simulated annealing. *science* 220 (4598), 671–680. doi:10.1126/science.220.4598.671
- Krauss, P., Prebeck, K., Schilling, A., and Metzner, C. (2019a). Recurrence resonance in three-neuron motifs. *Front. Comput. Neurosci.* 13, 64. doi:10.3389/fncom.2019.00064
- Krauss, P., Schuster, M., Dietrich, V., Schilling, A., Schulze, H., and Metzner, C. (2019c). Weight statistics controls dynamics in recurrent neural networks. *PLoS one* 14 (4), e0214541. doi:10.1371/journal.pone.0214541
- Krauss, P., Tziridis, K., Metzner, C., Schilling, A., Hoppe, U., and Schulze, H. (2016). Stochastic resonance controlled upregulation of internal noise after hearing loss as a putative cause of tinnitus-related neuronal hyperactivity. *Front. Neurosci.* 10, 597. doi:10.3389/fnins.2016.00597
- Krauss, P., Zankl, A., Schilling, A., Schulze, H., and Metzner, C. (2019b). Analysis of structure and dynamics in three-neuron motifs. *Front. Comput. Neurosci.* 13 (5), 5. doi:10.3389/fncom.2019.00005
- Langton, C. G. (1990). Computation at the edge of chaos: phase transitions and emergent computation. *Phys. D. Nonlinear Phenom.* 42 (1–3), 12–37. doi:10.1016/0167-2789(90)90064-v
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *nature* 521 (7553), 436–444. doi:10.1038/nature14539
- Legenstein, R., and Maass, W. (2007). Edge of chaos and prediction of computational performance for neural circuit models. *Neural Netw.* 20 (3), 323–334. doi:10.1016/j.neunet.2007.04.017
- Luczak, A., Barthó, P., and Harris, K. D. (2009). Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron* 62 (3), 413–425. doi:10.1016/j.neuron.2009.03.014
- Maheswaranathan, N., Williams, A. H., Golub, M. D., Ganguli, S., and Sussillo, D. (2019). Universality and individuality in neural dynamics across large populations of recurrent networks. *Adv. neural Inf. Process. Syst.* 2019, 15629–15641.
- Marder, E., and Bucher, D. (2001). Central pattern generators and the control of rhythmic movements. *Curr. Biol.* 11 (23), R986–R996. doi:10.1016/s0960-9822(01)00581-4
- Martínez, L. (2014). Introducing the levinthal's protein folding paradox and its solution. *J. Chem. Educ.* 91 (11), 1918–1923. doi:10.1021/ed300302h
- Marzola, P., Melzer, T., Pavesi, E., Gil-Mohapel, J., and Brocardo, P. S. (2023). Exploring the role of neuroplasticity in development, aging, and neurodegeneration. *Brain Sci.* 13 (12), 1610. doi:10.3390/brainsci13121610
- Maximilian Schäfer, A., and Zimmermann, H. G. (2006). "Recurrent neural networks are universal approximators," in *International conference on artificial neural networks* (Springer), 632–640.
- McDonnell, M. D., and Ward, L. M. (2011). The benefits of noise in neural systems: bridging theory and experiment. *Nat. Rev. Neurosci.* 12 (7), 415–425. doi:10.1038/nrn3061
- Metzner, C., and Krauss, P. (2022). Dynamics and information import in recurrent neural networks. *Front. Comput. Neurosci.* 16, 876315. doi:10.3389/fncom.2022.876315
- Metzner, C., Yamakou, M. E., Voelkl, D., Schilling, A., and Krauss, P. (2024). Quantifying and maximizing the information flux in recurrent neural networks. *Neural Comput.* 36 (3), 351–384. doi:10.1162/neco_a_01651
- Molgedey, L., Schuchhardt, J., and Schuster, H. G. (1992). Suppressing chaos in neural networks by noise. *Phys. Rev. Lett.* 69 (26), 3717–3719. doi:10.1103/physrevlett.69.3717
- Moss, F., Bulsara, A., and Shlesinger, M. F. (1993). Stochastic resonance in physics and biology. In *Proceedings of the NATO Advanced Research Workshop*. Karlsruhe, Germany. *J. Stat. Phys.*, 70.
- Moss, F., Ward, L. M., and Sannita, W. G. (2004). Stochastic resonance and sensory information processing: a tutorial and review of application. *Clin. Neurophysiol.* 115 (2), 267–281. doi:10.1016/j.clinph.2003.09.014
- Narang, S., Elsen, E., Diamos, G., and Sengupta, S. (2017). Exploring sparsity in recurrent neural networks. *arXiv Prepr. arXiv:1704.05119*.
- Natschläger, T., Bertschinger, N., and Legenstein, R. (2005). At the edge of chaos: real-time computations and self-organized criticality in recurrent neural networks. *Adv. neural Inf. Process. Syst.* 17, 145–152.
- Panagiotaropoulos, T. I., Kapoor, V., Logothetis, N. K., and Deco, G. (2013). A common neurodynamical mechanism could mediate externally induced and intrinsically generated transitions in visual awareness. *PLoS One* 8 (1), e53833. doi:10.1371/journal.pone.0053833
- Rajan, K., Abbott, L. F., and Sompolinsky, H. (2010). Stimulus-dependent suppression of chaos in recurrent neural networks. *Phys. Rev. E* 82 (1), 011903. doi:10.1103/physreve.82.011903
- Röder, K., Joseph, J. A., Husic, B. E., and Wales, D. J. (2019). Energy landscapes for proteins: from single funnels to multifunctional systems. *Adv. Theory Simulations* 2 (4), 1800175. doi:10.1002/adts.201800175
- Sandamirskaya, Y. (2013). Dynamic neural fields as a step toward cognitive neuromorphic architectures. *Front. Neurosci.* 7, 276. doi:10.3389/fnins.2013.00276
- Schilling, A., Gerum, R., Boehm, C., Rasheed, J., Metzner, C., Maier, A., et al. (2024). Deep learning based decoding of single local field potential events. *NeuroImage* 120696. doi:10.1016/j.neuroimage.2024.120696
- Schilling, A., Gerum, R., Metzner, C., Maier, A., and Krauss, P. (2022). Intrinsic noise improves speech recognition in a computational model of the auditory pathway. *Front. Neurosci.* 16, 908330. doi:10.3389/fnins.2022.908330
- Schilling, A., Sedley, W., Gerum, R., Metzner, C., Tziridis, K., Maier, A., et al. (2023). Predictive coding and stochastic resonance as fundamental principles of auditory phantom perception. *Brain* 146 (12), 4809–4825. doi:10.1093/brain/awad255
- Schilling, A., Tziridis, K., Schulze, H., and Krauss, P. (2021). The stochastic resonance model of auditory perception: a unified explanation of tinnitus development, zwick tone illusion, and residual inhibition. *Prog. Brain Res.* 262, 139–157. doi:10.1016/bs.pbr.2021.01.025
- Schrauwen, B., Buesing, L., and Legenstein, R. (2009). "On computational power and the order-chaos phase transition in reservoir computing," in *22nd annual conference on neural information processing systems (NIPS 2008)* (Vancouver, BC, Canada: NIPS Foundation), Vol. 21, 1425–1432.
- Schuecker, J., Goedeke, S., and Helias, M. (2018). Optimal sequence memory in driven random networks. *Phys. Rev. X* 8 (4), 041029. doi:10.1103/physrevx.8.041029
- Sietsma, J., and Dow, R. J. F. (1991). Creating artificial neural networks that generalize. *Neural Netw.* 4 (1), 67–79. doi:10.1016/0893-6080(91)90033-2
- Solé, R. V., and Miramontes, O. (1995). Information at the edge of chaos in fluid neural networks. *Phys. D. Nonlinear Phenom.* 80 (1–2), 171–180. doi:10.1016/0167-2789(94)00158-m
- Song, S., Sjöström, P. J., Reigl, M., Nelson, S., and Chklovskii, D. B. (2005). Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biol.* 3 (3), e68. doi:10.1371/journal.pbio.0030068
- Stein, R. B., Gossen, E. R., and Jones, K. E. (2005). Neuronal variability: noise or part of the signal? *Nat. Rev. Neurosci.* 6 (5), 389–397. doi:10.1038/nrn1668
- Toyoizumi, T., and Abbott, L. F. (2011). Beyond the edge of chaos: amplification and temporal integration by recurrent networks in the chaotic regime. *Phys. Rev. E* 84 (5), 051908. doi:10.1103/physreve.84.051908
- Van Laarhoven, P. J. M., Aarts, E. H. L., van Laarhoven, P. J. M., and Aarts, E. H. L. (1987). *Simulated annealing*. Springer.
- Wallace, E., Maci, H. R., and Latham, P. E. (2013). Randomly connected networks have short temporal memory. *Neural Comput.* 25 (6), 1408–1439. doi:10.1162/neco_a_00449
- Wang, X. R., Lizier, J. T., and Prokopenko, M. (2011). Fisher information at the edge of chaos in random boolean networks. *Artif. life* 17 (4), 315–329. doi:10.1162/artl_a_00041
- Ward, L. M. (2013). The thalamus: gateway to the mind. *WIREs Cognitive Sci.* 4 (6), 609–622. doi:10.1002/wcs.1256
- Watanabe, T., Lawson, R. P., Walldén, Y. S. E., and Rees, G. (2019). A neuroanatomical substrate linking perceptual stability to cognitive rigidity in autism. *J. Neurosci.* 39 (33), 6540–6554. doi:10.1523/jneurosci.2831-18.2019
- Weng, L. (2020). Exploration strategies in deep reinforcement learning. *Lilianweng. Github. io/lil-log*.
- Wolynes, P. G. (2015). Evolution, energy landscapes and the paradoxes of protein folding. *Biochimie* 119, 218–230. doi:10.1016/j.biochi.2014.12.007
- Wu, Y., and Koutstaal, W. (2020). Charting the contributions of cognitive flexibility to creativity: self-guided transitions as a process-based index of creativity-related adaptivity. *PLoS one* 15 (6), e0234473. doi:10.1371/journal.pone.0234473
- Zhang, C., Zhang, D., and Stepanyants, A. (2021). Noise in neurons and synapses enables reliable associative memory storage in local cortical circuits. *Eneuro* 8 (1), ENEURO.0302–20.2020. doi:10.1523/eneuro.0302-20.2020
- Zwanzig, R., Szabo, A., and Bagchi, B. (1992). Levinthal's paradox. *Proc. Natl. Acad. Sci.* 89 (1), 20–22. doi:10.1073/pnas.89.1.20