



Oscillations in Natural Neuronal Networks; An Epiphenomenon or a Fundamental Computational Mechanism?

Wolf Singer^{1,2,3} · Felix Effenberger¹

Received: 6 June 2024 / Revised: 4 February 2025 / Accepted: 7 February 2025
© The Author(s) 2025

Abstract

The first part of the paper is devoted to a comparison between the functional architectures of the cerebral cortex and artificial intelligent systems. While the two systems share numerous features, natural systems differ in at least four important aspects: i) the prevalence of recurrent connections, ii) the ability to use the temporal domain for computations, iii) the ability to perform "in memory" computations and iv) the prevalence of analog computations. The second part of the paper focuses on a simulation study that has been designed to answer the long-standing question of whether the oscillatory patterning of neuronal activity, which is a hallmark of natural systems, is an epiphenomenon of recurrent interactions or serves a functional role. To this end, recurrent neuronal networks were simulated to capture essential features of cortical networks, and their performance was tested on standard pattern recognition benchmark tests. In order to control the oscillatory regime of these networks, its nodes were configured as damped harmonic oscillators. By varying the damping factor, the nodes functioned either as leaky integrators or oscillators. It turned out that networks with oscillatory nodes substantially outperformed their non-oscillating counterparts. The reasons for this superior performance and similarities with natural neuronal networks are discussed. It is concluded that the oscillatory patterning of neuronal responses is functionally relevant because it allows the exploitation of the unique dynamics of coupled oscillators for analog computation.

Keywords Cerebral cortex · Oscillations · Synchrony · Computational neuroscience · Neuronal networks

✉ Wolf Singer
wolf.singer@brain.mpg.de

¹ Ernst Strüngmann Institute for Neuroscience in Cooperation with Max Planck Society, Frankfurt, Germany

² Max Planck Institute for Brain Research, Frankfurt, Germany

³ Frankfurt Institute for Advanced Studies, Frankfurt Am Main, Germany

Changing Views on Brain Dynamics

Metaphors

Ever since anatomists and physiologists have set out to analyze the functions of organs, attempts have been made to come up with analogies referring to contemporary, man-made artifacts: The heart a pump, the kidney a filter, and the liver a chemical plant. This also holds for the brain. However, finding plausible metaphors for the functions of the brain turned out to be less straightforward. The major problem is, of course, the notorious mind-body problem. So far, there is no example of a material system other than the brain that is capable not only of interacting with the material dimension of the world but also of generating immaterial realities such as the qualia of conscious experience, emotions, and beliefs. Yet, here we do not intend to discuss the question of whether the emergence of these immaterial phenomena can be accommodated in the framework of naturalistic explanations. Readers motivated to further explore this question are invited to consult a recent paper that contains abundant references to studies dealing with these problems (Singer, 2019).

However, certain functions such as performing logical operations and calculus, mediating sensory-motor reflexes, maintaining homeostasis, and programming complex motor responses turned out to be readily implementable in machines. Accordingly, the brain has been successively compared with ever more sophisticated technical devices. Nurtured by the monist manifesto of Julien Offray De La Mettrie, formulated in his book “L’homme machine” (1748), and by the behaviorist stance, which considered nervous systems as stimulus-response machines, the brain has been compared with mechanical calculators, hydro-pneumatic machines, electric circuits performing analog computations, cybernetic devices made up of interconnected black boxes accomplishing specific functions, and more recently digital computers performing complex logical operations, to name but a few. Influenced by the concepts of Gestalt psychology (Köhler, 1930) and theories on the dynamics of complex, self-organizing systems (Haken, 1977; Nicolis & Prigogine, 1977), a parallel line of more wholistic theories of brain dynamics was developed that emphasized analogies with complex dynamical systems (Breakspear, 2017). The core constituents of these theories are phenomena such as oscillations, propagating waves, interference, competition, and synergy. These phenomena are ubiquitous in organisms but also observed in inanimate systems. Examples are pattern-generating chemical diffusion-reaction systems or mechanical, optical, and electronic systems consisting of reciprocally coupled elements. Physical systems thought to capture some of the neuronal dynamics and imitate the respective computations are holograms, or more generally, any medium supporting propagation and interference of waves (Maksymov et al., 2022). The principle of holography uses the interference of waves, in this case, light waves, to encode information in a holistic, distributed manner. The interference pattern resulting from the encoding of a visual scene can be engraved in an optical medium, e.g., a glass plate. Upon illumination of any part of his plate with coherent light, the visual scene can then be reconstructed. Lately, even quantum systems are used as metaphors for brains as they share the brain’s abilities to solve hard optimization problems, such as the Traveling Salesman Problem and visual scene segmentation in a seemingly effortless way. In this case, the common denominator is that basic operations can be captured with the very general and powerful mathematical formulation of wave functions.

Artificial Systems

As this brief recapitulation of metaphors for the brain suggests, we are currently dealing with widely divergent classes of models, all of which capture some aspects of brain function. This makes it difficult to determine which concept fits best with our current knowledge of the structural and functional organization of natural brains.

At first sight, the answer seems straightforward. In the past two decades, progress in the design of computing algorithms, implementation of machine learning algorithms, and a tremendous increase in computing power have led to systems that outperform human brains in a number of functions that had been considered accomplishable only by highly evolved brains. We have already become used to the fact that mechanical and chemical sensors, electronic calculators, data mining systems, and autopilots outperform our cognitive abilities in some areas. However, realizing that there are now machines that excel in tasks that we consider as requiring reasoning, decision-making, educated judgments, and understanding of meaning is experienced as an attack on our human dignity. Generative large language models seemingly pass the Turing test and experienced doctors rely more and more on the diagnostic abilities of trained supercomputers. Are these systems then an appropriate model for our brains? Have we understood how our brains function, bypassing neurobiological investigations, by creating and, therefore, understanding systems that accomplish genuine brain functions? A number of arguments suggest that this might be the case. The architecture of the deep feed-forward neuronal networks (DNNs), which are currently mostly implemented on digital von Neumann architectures, shares features with natural neuronal networks. In fact, they are, in most cases, based on the perceptron model proposed in the fiftieth of the last century by Frank Rosenblatt (1958). This model was inspired by the seminal works of McCulloch and Pitts (1943) and Donald Hebb (1949), and aimed at capturing some of the wiring principles of natural neuronal networks (Fig. 1).




Fig. 1 Perceptron and feed-forward architectures of artificial neuronal networks. **A.** Schematic of the Perceptron architecture in which each node computes the weighted sum of its inputs, followed by the application of a non-linearity f . **B.** Schematic of a feed-forward network with an input layer, two hidden layers, and an output layer. Each circle indicates a Perceptron node as in **A**. Nodes code for features by means of convergent weight configurations (labelled line code)

In such architectures, the nodes of the network are capable of integrating converging input signals and converting them into a single output according to a pre-specified function and thus share rudimentary properties with neurons. Input signals are relayed in a serial way through feedforward connections via so-called hidden layers to an output layer. The feed-forward connections are convergent and distribute the signals of individual units of one layer onto nodes in the next layer, each node integrating signals from a large number of weighted inputs. Thus, inputs are transformed into distributed activation patterns in the hidden layers. By adjusting the weights of the feedforward connections, it becomes possible to improve the segregation of activation patterns in the hidden layer and ultimately enable the classification of these patterns with linear classifier units. It suffices to adjust inputs from the hidden layer nodes to the output layer so that a particular pattern in the hidden layer activates a corresponding unit in the output layer (the so-called read-out layer) preferentially. Even the latest deep learning architectures, such as convolutional or transformer networks, are constructed according to this principle. Temporal dynamics are avoided by design in order to facilitate highly parallelizable training.

It is worth mentioning that the main drivers of progress in the last decades have been technological rather than conceptual advances, such as the ever-growing availability and affordability of computational resources, in particular of GPUs allowing for highly parallel processing. These advances enabled the construction and training of larger and larger models on larger and larger data sets. Transformer networks, for example, were motivated by their ability to capture certain aspects of recurrent dynamics but still rely on feed-forward architectures that are easier to control. All these networks are trained with a method that is known as back-propagation of errors (or back-propagation, for short) (Linnainmaa, 1970; Werbos, 1982). Next to the increase in computing power, it was the refinement of this training algorithm that contributed essentially to the excellent performance of these networks.

In contrast to feed-forward architectures, training *recurrent* artificial networks on temporally organized inputs with the back-propagation algorithm has remained a challenge (Pascanu et al., 2013), despite important developments such as the introduction of gated units that simulate memory functions (Hochreiter & Schmidhuber, 1997). As a result, feed-forward architectures characterized by their basic principle of serial feed-forward processing still continue to dominate the field of machine learning. These systems learn essentially by establishing relations between a large set of input patterns and the activity of corresponding read-out units using labeled-line codes. Unlike natural systems they cannot exploit the temporal domain for their computation.

Natural Systems

These features of DNNs sharply contrast with those of natural neuronal networks, particularly those of the cerebral cortex and hippocampus. Although the processing streams in the cerebral cortex exhibit similarities to the serial organization of multi-layered DNNs, notable differences are evident. In the cortex, nodes within each layer interact reciprocally through abundant recurrent connections, and there are extensive feedback connections from higher to lower layers of the processing hierarchy. Due to this recurrent connectivity and the tendency of the neurons (network nodes) to oscillate, cortical networks display highly complex, non-linear dynamics both at rest and when stimulated. At the population level, these dynamics manifest as an ensemble of oscillations across a wide range of frequencies, extending from far below 1 Hz up to several hundred Hz. Furthermore, the nodes of the network conduct analog, non-linear computational operations on time-varying input

signals. These operations are primarily mediated by discrete frequency-modulated action potentials that reflect fluctuations in the sender's membrane potential, which are then converted by chemical synapses into graded fluctuations of the target neurons' membrane potential.

To convey an intuition about the properties of such recurrent networks and their potential to perform useful computations, the metaphor of a liquid medium is often put forward. Imagine a pond of water into which stones are thrown in a sequence at different locations (Fig. 2). Each impact will generate a spreading wave whose origin, amplitude and wavelength reflects the site, and the size, as well as the strength of the impact. After some time, the various waves will meet and generate a complex, high dimensional interference pattern that eventually fades. If one places several sensors into the pond that measure the amplitude, frequency, and phase of the local oscillations, it is, in most cases, possible to reconstruct when and where the various impacts have occurred. The water "remembers" the impacts as long as the waves persist, and the interference pattern contains all the information required to reconstruct the series of events (Fernando & Sojakka, 2003; Lu et al., 2020). Thus, the water performs several interesting computations. It transforms a stimulus into a stimulus-specific oscillation, distributes information about the spatial and temporal properties of the stimulus over the whole medium, establishes relations between the spatial and temporal parameters of different events through interference, and transforms the low-dimensional sequence of stimuli into a high-dimensional dynamic pattern that is at each time-point defined by the spatial distribution of the amplitude, the frequency and the phase of the oscillations. Due to fading memory and the high dimensionality of the dynamic landscape, the medium permits the superposition of waves and the simultaneous representation of information about temporally segregated events. These powerful computational primitives can be exploited by recurrent architectures that support the propagation of waves.

It is important to note that these computations are analog, not digital and that they are performed in the same medium that also serves as a memory. Thus, there is no separation into memory and processing circuits as in conventional von Neumann digital computers. The low-dimensional input pattern is converted into a very high-dimensional interference




Fig. 2 Storage, superposition and processing of information in high-dimensional dynamic state spaces using waves. Left: Water drops eliciting oscillations and traveling waves on a water surface that lead to dynamically evolving interference patterns. Right: Schematic of a recurrent neuronal network in which sensory and recurrent input result in similar dynamics such as traveling waves

pattern, exploiting constructive and destructive interference phenomena. Therefore, input patterns that overlap in a low-dimensional space and hence are difficult to distinguish from one another get well segregated in the high-dimensional state space of the interference patterns and hence can be classified more easily by appropriate read-out devices. Accordingly, this principle is exploited for numerous applications and is addressed as "liquid" or "reservoir" computing (Buonomano & Maass, 2009; Lukoševičius & Jaeger, 2009). As discussed above, perceptrons and DNNs also apply the principle of dimensionality expansion and reduction for pattern classification, but they do so by serial recombination of diverging and converging feed-forward connections rather than by exploiting dynamic features of network activity (Fig. 2).

Interestingly, the dynamics of recurrent neuronal networks resemble in several aspects those of the pond of water. Local activation of a node gives rise to spreading activity, so-called traveling waves, that inform with some latency other nodes in the network, and this reverberating activity keeps the trace of the stimulus alive for some time, a phenomenon addressed as "fading memory" (Nikolic et al., 2009). If the networks are perturbed by more than one stimulus delivered either at different sites or in a sequence, the resulting traveling waves will interact with one another and create a high-dimensional interference pattern. However, natural recurrent networks have additional properties that greatly enhance their computational capacity. Interactions in our water example are mostly homogeneous and restricted to nearest neighbours. By contrast, in biological neuronal networks, nodes are coupled through connections that can also mediate direct interactions between remote nodes, and, most importantly, these coupling connections are anisotropic and have temporally modulated properties (short and long-term plasticity). While genetic programs determine the basic layout of these connections, their detailed architecture is modified extensively by experience, resulting from interactions with the organism's *Umwelt*.

The Internal Model of the World

In the visual cortex, neurons responding to stimulus features that often co-occur in natural environments will be activated more often at the same time than neurons responding to features that are rarely contingent. The consequence is that these neurons will become more strongly connected with one another than neurons that respond to features that rarely co-occur (Hebb, 1949). The reason is that the reciprocal connections between neurons are endowed with correlation-sensitive synapses, so-called Hebbian synapses. "Neurons wire together if they fire together" (Löwel & Singer, 1992). In this way, statistical regularities of the environment get incorporated into the architecture of the dense network of recurrent cortical connections. This allows the brain to build an internal model of the world, formed by a series of so-called priors, by evaluating temporal correlations among neuronal discharges and translating these correlations into the functional architecture of recurrent connections. During early development, these activity-dependent modifications lead to macroscopic changes in connection architectures because a large fraction of initially formed connections is eliminated. Connections between neurons that are rarely activated together get pruned away. These activity-dependent modifications of recurrent coupling connections continue into adulthood and follow the same correlation rule but then are, in principle, confined to mostly reversible increases and decreases of synaptic efficiency (Artola et al., 1990). The formation of new long-range connections is no longer possible. These activity-dependent modifications of coupling are considered the mechanism by which the brain acquires and stores knowledge about the structure of the world, knowledge that is required

to interpret sensory signals (Rao & Ballard, 1999). Thus, the recurrent networks of natural brains have a number of interesting properties, and this raises the question of to which extent these are actually exploited for information processing.

Dynamic Encoding of Relations in Recurrent Networks

In Perceptron-based feed-forward architectures, relations among features in input patterns are encoded by converging connections onto conjunction-specific nodes (neurons), the responses of these nodes signaling the presence of a particular relation. The problem with this explicit representation of relations, also addressed as labeled line coding, is the combinatorial explosion of required conjunction-specific nodes as the number of relations increases. This makes the labeled line strategy, as well as its recent extension to transformer networks, costly in terms of hardware. A complementary strategy to capture relations among components or features relies on dynamic combinatorial codes, similar to those used by natural languages. This strategy was proposed by Donald Hebb (1949) more than half a century ago. It posits that relations among components should be encoded by transiently binding feature-selective neurons into functionally coherent assemblies, the Hebbian assembly. Through cooperative interactions, these neurons would collectively signal the presence of a particular constellation of components. The formation of such transient assemblies requires self-organized cooperativity among network nodes and is, therefore, difficult to implement in feed-forward architectures that cannot exploit time for coding relations. By contrast, the required interactions can be realized elegantly in recurrent network architectures, which incorporate time as a crucial coding dimension. According to Hebb's proposal, the recurrent, reciprocal connections must be endowed with correlation-sensitive synaptic plasticity mechanisms (Hebbian synapses) in order to preferentially stabilize the activity of assemblies of cells representing frequent or meaningful constellations of features. With exposure to natural visual scenes, nodes tuned to frequently cooccurring features would become coupled more strongly and, therefore, engage more readily in cooperative interactions when activated by the respective feature constellation. As discussed above, these important prerequisites have meanwhile been confirmed experimentally. In this coding scheme, conjunctions of features are represented by groups of temporarily cooperating nodes rather than individual conjunction-specific neurons. This allows the flexible and dynamical recombination of nodes to capture different constellations of features, elegantly overcoming the problem of the combinatorial explosion caused by labeled line codes. In artificial systems relying on labeled line codes, this problem is solvable only with brute force (hardware).

Another advantage of recurrent processing is that it can exploit not only the rate of discharges but also their relative timing for computations. This expands considerably the coding space and also accelerates processing speed. Donald Hebb (1949) had initially proposed that nodes forming a cooperating assembly should be distinguished by joint increases of discharge (firing) rate. And there is indeed evidence that recurrent interactions can enhance discharge frequency (Peron et al., 2020). However, the notion that cooperating neurons should be distinguished solely by joint increases in discharge rate has been challenged (for a review, see Singer, 1999). First, different simultaneously active and spatially intermingled assemblies are difficult to distinguish from one another if all neurons participating in assemblies simply discharge more vigorously—a complication addressed as the “superposition problem” (von der Malsburg & Schneider, 1986; Milner, 1992a, 1992b). Second,

increases in discharge rate are an ambiguous signature for a relational code because discharge rates also reflect stimulus energy and/or the quality of matches between stimulus and receptive field properties that need not necessarily serve feature binding. Third, distinguishing cooperating neurons on the mere basis of enhanced discharge rates slows down processing speed because discharge rates of cortical neurons are low and can carry only little information when integrated over short time intervals (Averbeck et al., 2006; Peron et al., 2020).

The serendipitous finding that neurons in the visual cortex can engage in synchronous oscillatory firing when driven by stimuli that possess related features offered solutions to the problems associated with rate-coded assemblies. Neurons in the visual cortex activated by continuous contours were found to synchronize their spike discharges with millisecond precision if located in functional columns segregated by up to 7 mm (Gray & Singer, 1989; Gray et al., 1989). Because synchronization enhances the impact of discharges in downstream targets (Abeles, 1991; Bruno & Sakmann, 2006; Salinas & Sejnowski, 2001), it was proposed that the joint increase in the salience of precisely synchronized discharges, rather than joint rate increases, would identify the transiently cooperating nodes of an assembly. This temporal code would substantially reduce the superposition problem because coincidence-sensitive downstream neurons can distinguish between different synchronous events with high temporal resolution (Singer, 1999; Milner, 1992a, 1992b; von der Malsburg & Buhmann, 1992; for reviews, see Singer & Gray, 1995). For the same reason, encoding relations by coincident firing rather than joint rate increases would allow for much faster detection of cooperating nodes (Van Rullen et al., 2005), and also render the signature of relatedness independent of rate fluctuations.

The Relation Between Gestalt Rules and Response Synchronization

Analysis of the synchronization phenomena observed in the visual cortex revealed that neurons transiently synchronized their discharges with millisecond precision when co-activated by continuous contours, collinearly aligned contour segments, contours sharing the same orientation, and contours moving with the same speed in the same direction (Gray et al., 1989; Singer & Gray, 1995). These feature constellations correspond to the basic Gestalt criteria for perceptual grouping, contiguity, continuity, similarity and common fate. This indicates that synchronization probability reflects the statistical regularities of natural environments and suggests that it could serve feature binding by generating Hebbian assemblies (for a review, see Singer, 1999). Support for this possibility came from the evidence that grouping by synchronization turned out to be dynamic and context-sensitive: Typically, when a single elongated moving contour is presented, all nodes activated by this contour synchronize their discharges (Gray et al., 1989; Livingstone, 1996). However, when two contours with different orientations overlap in space and move in different directions, the activated neurons split up into two groups (Castelo-Branco et al., 2000; Engel et al., 1991a; Kreiter & Singer, 1996). The neurons within each group discharge in synchrony, but there is no correlation at millisecond time scales among the discharges of neurons belonging to different groups. This self-organized grouping depends on the respective preferences of the neurons for the two contours. Neurons preferring the orientation of stimulus one over that of stimulus two join in the synchronization group driven best by stimulus one and vice versa. Thus, the synchronized firing of neurons is context-sensitive and signals whether the neurons are activated by a single “object” or two different “objects.”

As the discharge rates of all neurons were elevated to the same extent, this segregation of neurons into two distinct groups would not have been detectable by decoding only the rate responses. Simulations of recurrent networks with enhanced coupling of nodes tuned to groupable features have reproduced such context-dependent synchronization phenomena (König & Schillen, 1991; Schillen & König, 1991, 1994). The results of these and related studies eventually led to the proposal that transient and context-sensitive synchronization of discharges could serve as a mechanism for perceptual grouping and inspired the “binding by synchrony” (BBS) hypothesis (Singer, 1993).

Synchrony and Oscillations

Stimuli whose constellation of features matches the priors stored in the weight distribution of the recurrent connections not only induce synchronous firing among neurons tuned to features predicted to be related but also cause an oscillatory modulation of the responses in the gamma frequency range that is centered around 40 Hz. This frequency range is characteristic of local synchronization events, while synchronization at lower frequencies is observed over longer distances and is supposed to ensure the binding of processes between areas of the cerebral cortex. Stimuli rich in features matching stored priors, such as drifting gratings evoke synchronous oscillations in large ensembles of neurons, and these oscillations can then be detected even on the cortical surface or with extracranial recordings. Because increases in oscillatory power and spike synchronization covary, the important distinction between oscillatory patterning and spike synchronization got blurred, and the two phenomena are often regarded as equivalent (Stryker, 1989). Consequently, band-passed oscillations and their coherence rather than the timing relations among discharges of distributed neurons became the target variable in numerous studies investigating response synchronization (for review of the extensive literature on oscillations, see Fries, 2005, 2015; Buzsáki, 2006; Buzsáki et al., 2013; Singer, 2018). However, it has been argued that what matters for information processing are not the oscillations per se but the rate and the precise timing relations of discharges. The oscillatory patterning of responses so goes the argument, could have been simply an epiphenomenon of circuit interactions. Recurrent inhibition serves functions such as contrast enhancement, gain control, noise suppression, or winner-take-all mechanisms but inevitably also causes oscillations. Thus, it remained a matter of vivid debate whether oscillatory mechanisms play a role in cortical computations. Because oscillations are only sustained and stable in frequency under special stimulation conditions (Chen et al., 2014), vary in frequency as a function of stimulus energy and visual field eccentricity (Ray & Maunsell, 2010), and, if synchronized globally, limit the capacity of information transfer, oscillation-based synchrony was considered an epiphenomenon and irrelevant for information processing (for a review, see Ray & Maunsell, 2015). Moreover, more recent investigations in awake animals have revealed that synchronization of discharges, while often associated with oscillations, is not dependent on the occurrence of sustained oscillatory activity. Typically, episodes of synchronized firing are transient, occur in short bouts, follow at irregular intervals, and, if associated with oscillations, the oscillations tend to persist only over a few cycles (Chauvière & Singer, 2019; Lowet et al., 2016, 2017; Lundqvist et al., 2016). However, if nodes of neuronal networks have the propensity to oscillate - for which there is ample evidence (Börgers & Kopell, 2003; Buzsáki & Draguhn, 2004; Gray & McCormick, 1996; Jansen & Rit, 1995; Onorato et al., 2020; Spyropoulos et al., 2022) a host of dynamic phenomena emerge beyond the

occurrence of local oscillations, which may be relevant for computations. In simultaneous recordings from multiple nodes of cortical networks these dynamics manifest in long-range synchronization of discharges (Engel et al., 1991b), as frequency-varying oscillations (Kayser et al., 2003), transient synchronization or desynchronization of discharges (Glasser et al., 2016), resonance (Gross et al., 1972), entrainment (Tsao, et al., 2006), phase shifts (Hirabayashi et al., 2013), and traveling waves (Bell & Sejnowski, 1997; Olshausen & Field, 1996; Quiroga et al., 2005).

Unfortunately, it is notoriously difficult to determine in physiological experiments whether any of these dynamical phenomena serves a functional role because of the complexity, non-stationarity, and variability of the dynamics, the technical limitations to monitoring the activity of a sufficient number of nodes with the required temporal resolution, and the difficulty to apply knockout or gain-of-function strategies to prove causal relations. Interfering with single variables of such complex, highly integrated systems as recurrent networks inevitably impacts the behavior of the entire system, making it difficult to assess the putative function of a particular feature. Therefore, most of the evidence for a relation between spike synchronization, oscillatory patterning of neuronal discharges, and particular cognitive functions is only correlative in nature. So far, no experimental approach has been able to provide causal evidence for the functional relevance of oscillatory activity. This is a burning issue because the oscillatory patterning of neuronal responses is a ubiquitous phenomenon, and measures derived from these temporal patterns, such as frequency spectra, coherence, phase stability, and synchronization, are widely used as diagnostic tools of brain states in health and disease.

A Comprehensive Simulation Study

To alleviate the epistemic problem of distinguishing between relevant and epiphenomenal features of brain dynamics, we opted for a synthetic bottom-up approach. We simulated recurrent neuronal networks (RNNs) and implemented step-by-step characteristic features of neuronal circuits that had been identified experimentally. Subsequently, we tested the functional consequences of these additions by comparing the networks' performance on standard benchmark tests for pattern classification (Fig. 3).

In the first step, inspired by physiological experiments (Spyropoulos et al., 2022), we examined the consequences of configuring the network nodes as damped harmonic




Fig. 3 Harmonic Oscillator Recurrent Network (HORN) architecture. **A.** A HORN network consisting of recurrently coupled damped harmonic oscillator (DHO) nodes. Each node models the aggregate recurrent dynamics of a local cortical microcircuit. **B.** Relaxation dynamics of a DHO node for varying values of the natural frequency parameter ω (top) and the damping parameter γ (bottom)

oscillators (DHOs), a characteristic property of cortical microcircuits. We call such networks harmonic oscillator RNNs (HORNs).

To our surprise, we found that this was a decisive step. The HORNs outperformed, sometimes by orders of magnitude, RNNs without oscillating units such as, e.g., leaky integrators, with respect to parameter efficiency, task performance, learning speed, and noise tolerance. Encouraged by these findings, we continued to implement additional characteristics of cortical networks in HORNs. We endowed them with nodes that prefer different oscillation frequencies, introduced heterogeneous conduction velocities for the recurrent connections, and implemented Hebbian synapses. In addition, we recurrently coupled two HORNs in order to examine the putative advantages of multilayer architectures as found in the brain. The inclusion of these biologically inspired features resulted in further improved task performance without increasing the number of trainable variables (for details see Effenberger et al., 2025).

In-depth analyses of network dynamics uncovered a genuinely novel and powerful computational principle that exploits the rich transient and relaxation dynamics of coupled oscillators and casts a new view on recurrent network dynamics. Because the nodes of HORNs convert all stimuli in oscillatory responses, this allows the network to engage in highly dynamic interactions that are characterized by synchronization, resonance, entrainment, frequency and phase shifts, and dynamic gain modulation. Consequently, the networks can fully exploit both the spatial *and* the temporal dimensions for computations.

We consider it of particular importance that our network transforms all stimulus parameters into the continuous variables of oscillations because this permits analog computations. The traveling waves of reverberating activity and the resulting interference between the responses of the interconnected nodes permit a virtually simultaneous evaluation of both spatial *and* temporal relations between multiple, spatially and temporally segregated stimuli. And because the functional architecture of the recurrent coupling connections represents an internal model of the world, this computation within memory allows for a highly parallelized match of sensory evidence with stored priors. All these operations are crucial prerequisites for perceptual processes such as feature binding, scene segmentation, and, ultimately, the representation of complex perceptual objects.

It is noteworthy to state that in this framework, the representations of perceptual objects do not consist of pattern-specific read-out units but are distributed across nodes, non-stationary and holistic. Representations consist of dynamic landscapes that exhibit a complex and fine-grained correlation structure (Fig. 4). As all nodes, not only those stimulated directly, contribute to these representations, stimulus-specific information can be extracted from all nodes, like in a hologram. In the cerebral cortex, the read-out of these activity landscapes does not require convergence of the nodes' activity onto classifier units. The distributed activity patterns serve, in turn, as highly parallel input to downstream recurrent neuronal networks whose nodes are reciprocally connected to nodes of the respective upstream areas through divergent and convergent feed-forward and feed-back connections. Thus, there is no bottleneck in the processing stream from sensory to executive cortical areas. Rather, distributed dynamic representations emerge at all stages of processing, including the motor-cortex. A motor command needs to orchestrate the cooperation of myriads of muscle cells and is best implemented by a distributed spatially and temporally structured activity landscape that is generated according to the same principles as the representations in sensory cortices.

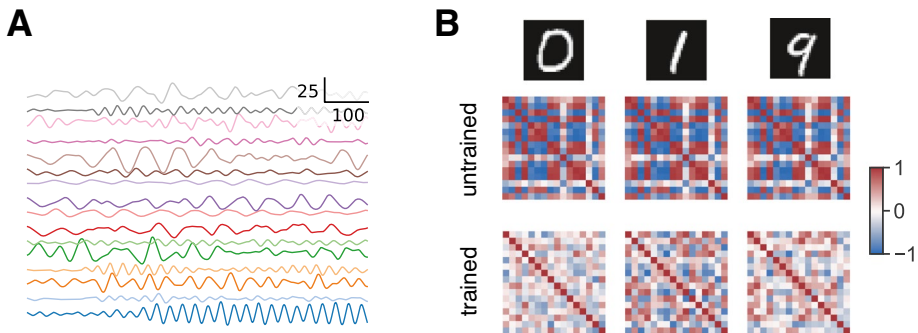


Fig. 4 Dynamics and correlation structure in HORN networks. **A.** Dynamics of 16 DHO nodes of a heterogeneous HORN trained on MNIST digit recognition resulting from the stimulation with a serialized MNIST digit. Vertical axis shows unit amplitude, horizontal axis shows time. **B.** Pairwise Pearson correlation coefficients of nodal dynamics of a 16 node HORN before (top row) and after learning (bottom row) show the emergence of learning-induced, fine-scale, stimulus-specific correlation structures in HORN network dynamics. Columns show different MNIST digit classes. Value of correlation coefficient color-coded (scale on right)

Documentation of HORNs Basic Structural and Functional Features

In the following, a few of the most characteristic properties of the HORNs will be presented to support the conclusions summarized above.

Oscillating Network Nodes

To facilitate the analysis of principles, we designed a simple RNN and configured the nodes as damped harmonic oscillators (DHOs). We opted for DHOs for the following reasons: First, they are the canonical implementation of an oscillatory process. Second, in neuronal microcircuits, damped harmonic oscillations generically result from excitatory-inhibitory interactions or negative feedback subject to a damping force (Buzsáki & Draguhn, 2004; Jansen & Rit, 1995). Moreover, their dynamics are easily interpretable and controllable by two parameters. Each DHO node has one state variable x , the oscillator's time-varying amplitude, and three parameters that jointly control its relaxation dynamics: the natural frequency ω , the damping coefficient γ , and an excitability coefficient α . This makes DHOs the most basic choice for introducing oscillations at the level of a single network node (Fig. 5). It is important to note that already a single DHO node possesses intriguing computational capabilities that are not accessible to non-oscillating nodes such as leaky integrators. DHOs can gain-modulate an input signal in a non-linear way as a function of its frequency profile and encode stimulus information in the oscillation phase. Moreover, the combination of the input non-linearity with the biologically inspired feedback connections endows the DHO nodes with a dynamical repertoire that goes far beyond that of a classical damped harmonic oscillator. Among other things, it allows nodes to express self-sustained oscillations, and to resonate at fractional harmonics of their natural frequency.

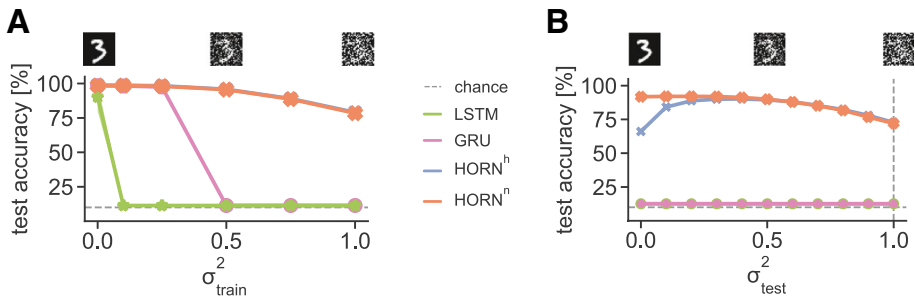


Fig. 5 HORN learning and inference on sMNIST corrupted by varying levels of additive white Gaussian noise. **A.** Maximal test accuracy of different RNN architectures (10^4 trainable parameters) after training on noisy sMNIST for 100 epochs as a function of training stimulus noise level σ_{train}^2 . **B.** Test accuracy of different network architectures trained on noisy sMNIST with $\sigma_{\text{train}}^2 = 1$ (dashed vertical line) as a function of test stimulus noise level σ_{test}^2 . Lines show means calculated over 1000 test stimuli

Homogeneous Networks

We first designed homogeneous HORN (HORN^h) networks with all-to-all connectivity, no conduction delays, and identical parameter values ω , γ , α , for all nodes. As cortical circuits operate in a balanced state where excitatory and inhibitory drive cancels in the mean (E/I balance), emphasizing the significance of fluctuations, we chose to couple DHO units on their velocity term. Stimuli were presented to the networks in the form of time series. We tested the networks on several well-established pattern recognition tasks, one of those is the MNIST handwritten digit recognition dataset. To transform the image of an MNIST digit into a time series, henceforth addressed as an sMNIST stimulus, intensity values were collected in scan-line order from top left to bottom right of the images. The networks were initially trained for the classification of sMNIST stimuli in a supervised way using the backpropagation through time (BPTT) algorithm. As DHOs resonate with input frequencies around ω_v , values of ω were chosen that enabled the nodes to extract and retain the information contained in time series resulting from continuous lines of different slopes. Thus, by adjusting ω , the network can be adapted to the statistical regularities of sMNIST stimuli, and this setting of priors enhances performance by allowing feature extraction through resonance rather than just through selective recombination of convergent input connections. To the best of our knowledge, this strategy of feature extraction has not been proposed previously and crucially depends on units or microcircuits with the propensity to oscillate.

Our benchmark tests revealed that these simple homogeneous HORN networks substantially outperformed state-of-the-art gated RNN architectures with respect to learning speed, absolute performance, and noise tolerance at the same number of learnable parameters. In addition, HORNs proved extraordinarily noise resilient, showing only a gradual decline in task performance with increasing noise levels, in contrast to other architectures studied (Fig. 6).

We identified several reasons for the high performance of the HORNs. (i) The kinetics of each DHO node turn arbitrary input into harmonic oscillations, thereby structuring the networks' state spaces and allowing for dynamic phenomena such as fading memory, resonance, entrainment, fine-scale synchronization, phase shifts, and desynchronization both on the single node and on the network level. These dynamic

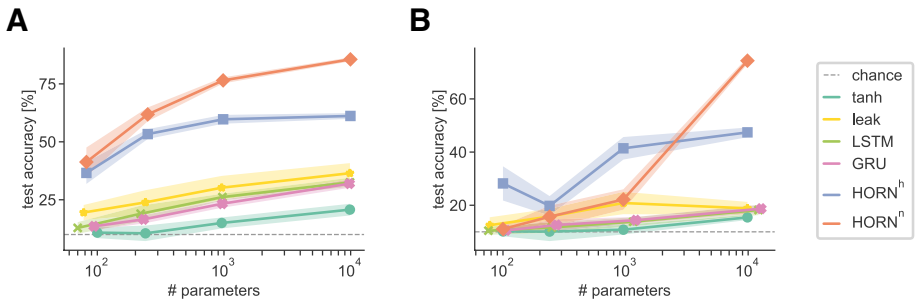


Fig. 6 Performance of different RNN architectures in classifying shuffled sequential MNIST (psMNIST) stimuli as a function of system size. Legend on the right. **A.** Test accuracy after 400 back-propagation learning steps. **B.** Test accuracy after 100 training epochs. Note the increased task performance of the heterogeneous HORN in comparison to the homogeneous network in the case of the spectrally more complex psMNIST stimuli

phenomena can then be exploited for learning. (ii) “Innate” preference for stimulus features (controlled by the values of ω, γ, α) allowed individual nodes to efficiently extract and encode stimulus features already in an untrained network by means of resonance. As expected, before learning, the dynamics of homogeneous HORN networks is dominated by large-scale synchronization among the nodes. As learning progresses, global synchronization decreases, which enhances the dimensionality of network dynamics. This reduction in global synchronization is accompanied by the emergence of complex, spatiotemporally structured correlation and higher-order synchronization patterns that are stimulus-specific, well segregated in the high-dimensional activity landscape of the network, and well classifiable by a linear read-out. For a detailed description of these properties, see Effenberger et al., 2025.

Heterogeneous Networks

The structural and functional organization of mature cortical networks is characterized by heterogeneity (Murray et al., 2014; Wang et al., 2006). However, it is unclear whether this variability plays a functional role. To test whether increasing network heterogeneity enhances performance, we simulated non-homogeneous HORNs (HORN^n) in which each node had a different natural frequency, damping coefficient, and excitability. As expected, heterogeneous HORNs responded already in the untrained state with more complex and less globally synchronized patterns. As in the homogeneous case, global synchrony decreased further as learning progressed. Here again, the decrease in global synchrony led to a significant increase in the dimensionality of the dynamics, but the dimensionality of the dynamics was much higher compared to their homogeneous counterparts already in the first training steps. A comparison between homogeneous and heterogeneous HORNs revealed the superior performance of the latter with respect to learning speed and noise tolerance, in particular for the more challenging psMNIST data set in which the pixel locations are randomized, resulting in more complex spectral properties of the stimuli (Fig. 7). We found that heterogeneity increases network expressivity and can alleviate the need for computationally expensive (and for biological systems likely unfeasible) parameter searches.




Fig. 7 Performance of HORNs with connection delays. Maximum test accuracy on psMNIST after 400 training steps (left) and after 4690 training steps (right), respectively, as a function of maximal synaptic delay d_{\max} . For each network, connection delays were sampled from a uniform distribution $U([1, d_{\max}])$. A network with $d_{\max}=1$ corresponds to a regular HORN. Lines show mean performance over 10 randomly initialized networks, shaded areas standard deviation

Delays

In the cerebral cortex, neurons interact through relatively slow conducting nerve fibers (0.5 to 10 m/s), which introduces considerable, widely scattered, distance-dependent coupling delays (Ferster & Lindström, 1983). To test the influence of introducing coupling delays on task performance, we started with a HORN^h and endowed all recurrent connections with uniformly distributed variable coupling delays (d). This manipulation increased HORN performance in both maximal classification accuracy and learning speed on psMNIST (Fig. 8). Increasing d_{\max} , which results in greater heterogeneity, was found to increase task performance, and this gain of function increased with increasing values of d_{\max} . Thus, as for the preferred oscillation frequencies, heterogeneity in conduction delays enables the generation of more diverse spatiotemporally structured activity landscapes in HORNs,




Fig. 8 Hebbian learning in HORNs. **A.** Scatter plots of connection weights W_{hh} and mean cross-correlation coefficients CC_{ij} of node activities of a 64-node homogeneous HORN before (left) and after training on psMNIST for 500 training steps (right). The CC_{ij} were computed for a random, but fixed set of 100 samples. Linear regression lines shown in red. Histograms on the top and right show marginal distributions of CC_{ij} and W_{hh} , respectively. Note the bimodal distribution of correlation coefficients with modes around -1 , 1 in the untrained state and the more decorrelated network activity resulting after learning. **B.** Performance of 64-node homogeneous HORNs as a function of training steps when instances of the same network are trained with correlation-based Hebbian (suffix Hebb+) or anti-Hebbian (suffix Hebb-) learning rules, compared to instances trained with BPTT and plastic connections W_{hh} (suffix BPTT) and with frozen W_{hh} (suffix reservoir). The input and readout parameters are trained with BPTT for all instances. Curves show mean performance over 10 network instances with random weight initialization, shaded areas standard deviation. Note the strong performance of the anti-Hebbian rule for this initially highly synchronized network

thereby increasing the dimensionality of the networks' state space. This again results in enhanced HORN performance and alleviates the need for parameter tuning, in particular for datasets with complex spectral properties or changing noise characteristics. Hence, we hypothesize that the heterogeneity of natural systems is likely not a reflection of unavoidable imprecision but developed in order to increase the dimension of the available coding space and to enable the networks to cope with a large variety of stimuli.

Hebbian Learning

Experimental evidence suggests that biological neuronal networks learn to represent characteristic feature constellations of perceptual objects through Hebbian modifications of synaptic gain (see above and Hebb, 1949). Thus, neurons that encode features that frequently co-occur become more strongly coupled and, therefore, self-organize to form a functionally coherent stimulus-specific assembly (Lazar et al., 2021).

To test whether gradient-based training of HORNs produces weight distributions compatible with those predicted by Hebbian learning principles, we compared recurrent weight distributions and the corresponding correlation structure of responses before and after learning. To our surprise, we found that the changes in synaptic weights of HORNs induced by BPTT are similar to those predicted by a Hebbian mechanism. The BPTT learning algorithm apparently capitalizes on the stimulus-specific correlation structure of network activity and enhances those connections that induce correlation patterns characteristic of a particular stimulus. As expected, heterogeneous HORNs could exploit stimulus-specific correlation structures right from the beginning of training. Homogeneous HORNs, by contrast, must first learn to desynchronize to increase the dimensionality of their state space, which slows down learning speed.

To investigate whether BPTT could be substituted by correlation-based learning, we implemented simple unsupervised additive (conventional) Hebbian, as well as anti-Hebbian rules for the activity-dependent modifications of the recurrent connections. During training, we disabled BPTT learning on the recurrent connections while keeping it active for the input and readout connections. As a performance baseline, we trained a HORN in which the recurrent connections were kept fixed, and only the input and readout connections were trained. In the case of homogeneous networks, the application of a fully unsupervised Hebbian rule resulted in a task performance surprisingly close to that attained when the recurrent connections were trained by the fully supervised BPTT rule (Fig. 9). In conclusion, the propensity of the network nodes to engage in oscillatory activity and the resulting amplification of specific correlations by *resonance* enabled the Hebbian mechanism to install stimulus-specific priors in the network architecture and to thereby orthogonalize representations in the high dimensional state space.

Overall, the results demonstrate that unsupervised Hebbian learning at the level of recurrent connections in HORNs aids in the segregation of stimulus-specific dynamic states, thereby facilitating their classification.

Spontaneous and Evoked Activity

Cortical networks are spontaneously active, and in sensory cortices, stimulation typically causes a reduction in the variability of this activity and leads to the emergence of stimulus-specific substates (Bányai et al., 2019; Berkes et al., 2011; Churchland et al.,

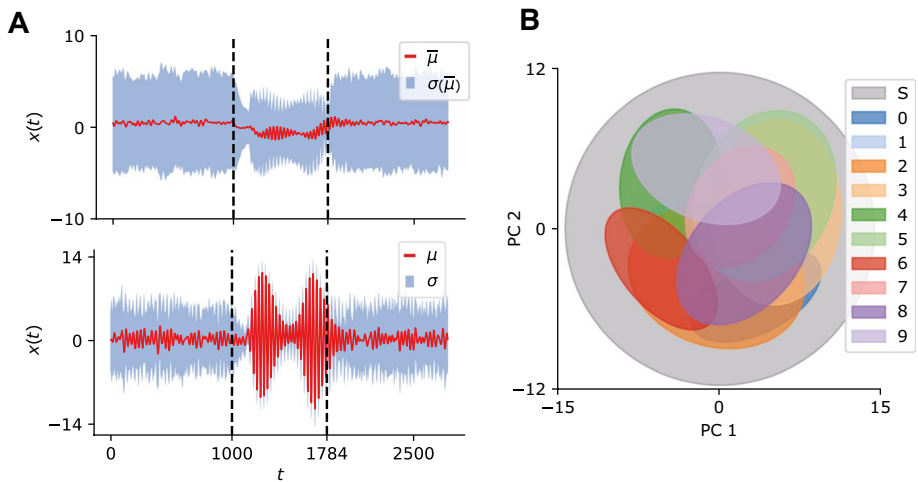


Fig. 9 Spontaneous and evoked activity in a 16-node heterogeneous HORN trained on sMNIST. **A.** Mean and standard deviation of amplitude dynamics computed over 500 samples of the sMNIST digit 0. Top: Mean amplitude μ and standard deviation of mean amplitude $\sigma(\mu)$ computed for the entire network. Bottom: Mean and standard deviation of amplitude dynamics of one selected DHO node. Dashed lines mark stimulus onset ($t=1000$) and offset ($t=1784$), respectively. **B.** Sliding window analysis of principal components of network activity (window size 20 time steps). Ellipses show 2σ confidence intervals, MNIST stimulus classes are color-coded (legend on right, S = spontaneous activity)

2010). To model spontaneous activity in HORNs, we subjected each DHO node to random discrete jumps in its velocity by a given amount according to a Poisson process.

After training a spontaneously active heterogeneous HORN on sMNIST, we analyzed network activity before, during, and after stimulation with sMNIST stimuli. During stimulation, spontaneous activity was replaced by the structured time series corresponding to the sMNIST digits. Analysis of the activity of both single nodes and the entire network showed that the variance of activity was high before stimulation, decreased during stimulation, and then recovered to pre-stimulus levels after stimulus offset. Principal component analysis (PCA) revealed that the state space of spontaneous activity spans a large but confined space that comprised the subspaces of stimulation-induced stimulus-specific response vectors. The reason for this confinement of the spontaneous state is that the connectivity of the trained HORN is not random but highly structured through learning. In trained HORNs, the weight distribution of the recurrent connections reflects the statistical regularities of the learned sMNIST stimuli. This structure of the coupling connections, in turn, shapes the dynamics of the network. Thus, spontaneous activity can be considered as a superposition of fragments of learned stimulus-specific representations. The temporal evolution of network dynamics in PCA space shows that network dynamics rapidly converges to a stimulus-specific substate once a stimulus is presented and that this specificity fades after stimulus offset (Fig. 9). At the same time, this collapse towards a stimulus-specific substate is associated with a reduction of variance of activity. Again, very similar, stimulation-dependent changes in network dynamics have been reported for natural networks (Churchland et al., 2010; Lazar et al., 2021; Singer, 2021).

Concluding Remarks

The implementation of characteristic features of the mammalian cerebral cortex in simulated RNNs uncovered, to the best of our knowledge, a genuinely novel and extremely powerful computational principle that exploits the unique dynamics of coupled oscillators and drastically enhances the performance of a recurrent network in a pattern classification task. Recurrent networks without oscillating nodes can, under certain circumstances, also engage in oscillatory activity. However, these oscillations are an emergent property of network interactions, highly variable and not controllable. Hence, this makes it difficult to determine their putative functions. We consider it unlikely that they can support computations and learning. In HORNs, the oscillations are tightly controlled and this allowed us to assess the functions of oscillations per se by comparing the performance of networks with and without oscillating nodes. The reason why oscillating nodes boost performance is that the network can take advantage of the rich dynamics of reciprocally coupled oscillators (Buzsáki & Wang, 2012; Cardin et al., 2009; Gray et al., 1989; Lachaux et al., 1999). These dynamics allow HORNs to exploit for computations not only spatial distributions of response amplitudes but also temporal variables such as resonance, entrainment, synchrony and asynchrony, reverberation, phase shifts, and fading memory. This also sets them apart from models based on attractor dynamics, such as the ones realized in Ising models and their associated Hopfield networks (Hopfield, 1982). Importantly, single network nodes in our model are intrinsically feature selective not only as a result of specific combinations of input connections, the strategy used in labeled line coding, but because of preferred resonance frequencies and non-linear gain modulation. Moreover, the oscillatory nature of the nodes' activity turns any input into a damped harmonic oscillation. Consequently, stimulus representations oscillate, a property that can be exploited for faster learning (see Dubinin & Effenberger, 2023 for a detailed analysis).

The DHO units in a HORN network collectively process stimuli in a fully distributed manner by converting sensory input into waves that spread and give rise to complex interference patterns (Hughes, 1995; Muller et al., 2018). This representation provides a coding space of massive dimensionality. It allows the networks to cope effectively with the encoding of time series and the superposition of information about spatially and temporally segregated events.

The introduction of physiologically plausible heterogeneity further enhanced performance at the same number of trainable parameters because it increased the dimensionality of the networks' state spaces and allowed them to simultaneously process stimuli on several time scales. Such heterogeneity endowed even untrained networks with sensitivity to a broader range of correlation structures that could be exploited by subsequent learning, accelerating learning speed. Most importantly, heterogeneity enhanced the ability to process stimuli with novel characteristics, alleviating the need for fine-tuning or the cumbersome search for optimal stimulus-specific parameters. Thus, heterogeneity is functionally beneficial and likely not a by-product of nature's imprecision.

Quite unexpectedly, gradient-based BPTT learning in HORNs resulted in changes in the coupling strengths of the recurrent connections that closely resemble those predicted by Hebbian principles. This provides strong support for the concept formulated by Singer (2021) that biologically plausible unsupervised correlation-based learning can install priors about statistical regularities of stimuli in the network's synaptic weight distributions, which, in turn, shape dynamic stimulus representations useful for stimulus discrimination. Our results provide a proof of concept that unsupervised learning rules can, in principle,

replace the physiologically implausible back-propagation method. Finally, spontaneously active HORNs were shown to reproduce important hallmarks of experimental findings, such as the stimulus-induced reduction of variance (Churchland et al., 2010) and the transient convergence of network dynamics to stimulus-specific substates (Singer, 2021).

Taken together, these findings not only unveil powerful novel principles based on analog computation but also allow for a functional interpretation of numerous experimentally verified physiological phenomena whose roles in information processing so far have been elusive or have caused controversial discussions. Plausible functional roles can now be assigned to (i) the propensity of nodes to oscillate and the resulting dynamical phenomena such as synchronization, desynchronization, resonance, entrainment, and traveling waves (Cardin et al., 2009; Davis et al., 2020; Doelling & Poeppel, 2015; Lachaux et al., 1999; Muller et al., 2018; Singer, 1993), (ii) the diversity of preferred oscillation frequencies, their non-stationarity and context dependence (Llinas, 1988; Engel et al., 1992), (iii) the heterogeneity of the conduction velocities of recurrent connections (Swadlow, 1985; Waxman, 1980), (iv) the Hebbian adaptivity of recurrent connections (Magee & Johnston, 1997; Markram et al., 1997), (v) the emergence of context-dependent dynamic receptive fields by network interactions (Blakemore et al., 1970; David, 2004), (vi) the decrease of oscillation frequencies in higher areas of the cortical processing hierarchy (Markov et al., 2013), and (vii) the reduction of variance in network dynamics during stimulus presentation (Churchland et al., 2010). The simulations also suggest a physiologically plausible scenario for the rapid and parallel matching of sensory evidence with stored priors through self-organized convergence of network dynamics to classifiable, stimulus-specific, dynamic substates. These substates consist of highly structured, high-dimensional dynamical landscapes that unfold in amplitude, frequency, and phase space.

In essence, such networks of coupled oscillators perform highly parallelized analog computations in very high-dimensional state space. This allows the networks to relate a large number of input variables simultaneously. We propose that this is why biological systems can solve extremely complex problems such as feature binding, scene segmentation or speech parsing so fast and efficiently despite the slowness of their constituting neurons. Problems of this type require the virtually simultaneous analysis of a vast number of nested relations and are notoriously difficult to solve with digital von Neumann computing architectures.

Exploiting quantum effects was one possibility for an efficient solution to such problems. Yet, it is commonly believed that biological systems cannot exploit quantum effects for computation. Therefore, we hypothesize that nature solves such hard problems employing analog computations of the kind described in this paper. At the conceptual level, the highly parallel, analog, and holistic computations realized in HORNs share numerous similarities with the computations performed in quantum systems. The superposition and interactions of phases in case of biological networks share similarities with the superposition and interaction of spins in physical systems (Moy et al., 2022). Being a non-quantum model, HORNs cannot emulate the phenomenon of quantum entanglement. However, HORNs and, most likely also, cortical networks exploit their ability to use the interference of superimposed waves for highly parallelized analog computations.

We think that it will be possible to implement the computational principle presented here in analog hardware that, in contrast to quantum systems, runs at room temperature, is miniaturisable, and highly energy efficient. We predict that this novel principle will likely enable the design of self-adapting and learning devices for machine learning applications that can ideally complement existing digital technologies.

Author Contribution W.S. and F.E. wrote the manuscript together. W.S. focused on the neurobiological and F.E. on the computational issues.

Funding Open Access funding enabled and organized by Projekt DEAL. Koselleck-Project, DFG (GZ: SI 505/22–1) and Ernst Strüngmann Institute for Neuroscience.

Data Availability No datasets were generated or analyzed during the current study.

Declarations

Ethical Approval Not applicable.

Competing Interests The authors declare no competing interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abeles, M. (1991). *Corticonics*. Cambridge University Press.
- Artola, A., Bröcher, S., & Singer, W. (1990). Different voltage-dependent thresholds for the induction of long-term depression and long-term potentiation in slices of the rat visual cortex. *Nature*, *347*, 69–72.
- Averbeck, B. B., Latham, P. E., & Pouget, A. (2006). Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, *7*, 358–366.
- Bányai, M., Lazar, A., Klein, L., Klon-Lipok, J., Stippinger, M., Singer, W., & Orbán, G. (2019). Stimulus complexity shapes response correlations in primary visual cortex. *Proceedings of the National Academy of Sciences of the USA*, *116*(7), 2723–2732.
- Bell, A. J., & Sejnowski, T. J. (1997). The “independent components” of natural scenes are edge filters. *Vision Research*, *37*(23), 3327–3338.
- Berkes, P., Orbán, G., Lengyel, M., & Fiser, J. (2011). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science*, *331*, 83–87.
- Blakemore, C., Nachmias, J., & Sutton, P. (1970). The perceived spatial frequency shift: Evidence for frequency-selective neurones in the human brain. *The Journal of Physiology*, *210*(3), 727–750.
- Börger, C., & Kopell, N. (2003). Synchronization in networks of excitatory and inhibitory neurons with sparse, random connectivity. *Neural Computation*, *15*, 509–538.
- Breakspear, M. (2017). Dynamic models of large-scale brain activity. *Nature Neuroscience*, *20*, 340–352.
- Bruno, R. M., & Sakmann, B. (2006). Cortex is driven by weak but synchronously active thalamocortical synapses. *Science*, *312*, 1622–1627.
- Buonomano, D. V., & Maass, W. (2009). State-dependent computations: Spatiotemporal processing in cortical networks. *Nature Reviews Neuroscience*, *10*, 113–125.
- Buzsáki, G. (2006). *Rhythms of the Brain*. Oxford University Press.
- Buzsáki, G., & Draguhn, A. (2004). Neuronal oscillations in cortical networks. *Science*, *304*, 1926–1929.
- Buzsáki, G., Logothetis, N., & Singer, W. (2013). Scaling brain size, keeping timing: Evolutionary preservation of brain rhythms. *Neuron*, *80*(3), 751–764.
- Buzsáki, G., & Wang, X.-J. (2012). Mechanisms of gamma oscillations. *Annual Review of Neuroscience*, *35*, 203–225.
- Castelo-Branco, M., Goebel, R., Neuenschwander, S., & Singer, W. (2000). Neural synchrony correlates with surface segregation rules. *Nature*, *405*, 685–689.

- Cardin, J. A., Carlén, M., Meletis, K., Knoblich, U., Zhang, F., Deisseroth, K., Tsai, L.-H., & Moore, C. I. (2009). Driving fast-spiking cells induces gamma rhythm and controls sensory responses. *Nature*, 459, 663–667.
- Chauvière, L., & Singer, W. (2019). Neurofeedback training of gamma oscillations in monkey primary visual cortex. *Cerebral Cortex*, 29(11), 4785–4802.
- Chen, M., Yan, Y., Gong, X., Gilbert, C. D., Liang, H., & Li, W. (2014). Incremental integration of global contours through interplay between visual cortical areas. *Neuron*, 82, 682–694.
- Churchland, M. M., Yu, B. M., Cunningham, J. P., Sugrue, L. P., Cohen, M. R., Corrado, G. S., Newsome, W. T., Clark, A. M., Hosseini, P., Scott, B. B., Bradley, D. C., Smith, M. A., Kohn, A., Movshon, J. A., Armstrong, K. M., Moore, T., Chang, S. W., Snyder, L. H., Lisberger, S. G., ... Shenoy, K. V. (2010). Stimulus onset quenches neural variability: A widespread cortical phenomenon. *Nature Neuroscience*, 13(3), 369–378.
- David, S. V. (2004). Natural stimulus statistics alter the receptive field structure of V1 neurons. *Journal of Neuroscience*, 24(31), 6991–7006.
- Davis, Z. W., Muller, L., Martínez-Trujillo, J., Sejnowski, T., & Reynolds, J. H. (2020). Spontaneous travelling cortical waves gate perception in behaving primates. *Nature*, 587, 432–436.
- Doelling, K. B., & Poeppel, D. (2015). Cortical entrainment to music and its modulation by expertise. *Proceedings of the National Academy of Sciences of the USA*, 112(45), E6233–E6242.
- Dubin, I., & Effenberger, F. (2023). Fading memory as inductive bias in residual recurrent networks. *arXiv*, July 28, 2023, 1–25. (<https://arxiv.org/abs/2307.14823>)
- Effenberger, F., Carvalho, P., Dubinin, I., & Singer, W. (2025). The functional role of oscillatory dynamics in neocortical circuits: A computational perspective. *Proceedings of the National Academy of Sciences of the USA*, 122(4), 1–12.
- Engel, A. K., König, P., Kreiter, A. K., Schillen, T. B., & Singer, W. (1992). Temporal coding in the visual cortex: New vistas on integration in the nervous system. *Trends in Neurosciences*, 15, 218–226.
- Engel, A. K., König, P., & Singer, W. (1991a). Direct physiological evidence for scene segmentation by temporal coding. *Proceedings of the National Academy of Sciences of the USA*, 88, 9136–9140.
- Engel, A. K., König, P., Kreiter, A. K., & Singer, W. (1991b). Interhemispheric synchronization of oscillatory neuronal responses in cat visual cortex. *Science*, 252, 1177–1179.
- Felleman, D. J., & van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex*, 1, 1–47.
- Fernando, C., & Sojakka, S. (2003). Pattern recognition in a bucket. *Lecture Notes in Computer Science*, 2801, 588–597.
- Ferster, D., & Lindström, S. (1983). An intracellular analysis of geniculate-cortical connectivity in area 17 of the cat. *The Journal of Physiology*, 342(1), 181–215.
- Fries, P. (2015). Rhythms for cognition: Communication through coherence. *Neuron*, 88, 220–235.
- Fries, P. (2005). A mechanism for cognitive dynamics: Neuronal communication through neuronal coherence. *Trends in Cognitive Sciences*, 9(10), 474–480.
- Glasser, M. F., Coalson, T. S., Robinson, E. C., Hacker, C. D., Harwell, J., Yacoub, E., Ugurbil, K., Anderson, J., Beckmann, C. F., Jenkinson, M., Smith, S. M., & Van Essen, D. C. (2016). A multi-modal parcellation of human cerebral cortex. *Nature*, 536, 171–178.
- Gray, C. M., König, P., Engel, A. K., & Singer, W. (1989). Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature*, 338, 334–337.
- Gray, C. M., & McCormick, D. A. (1996). Chattering cells: Superficial pyramidal neurons contributing to the generation of synchronous oscillations in the visual cortex. *Science*, 274, 109–113.
- Gray, C. M., & Singer, W. (1989). Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. *Proceedings of the National Academy of Sciences of the USA*, 86, 1698–1702.
- Gross, C. G., Rocha-Miranda, C. E., & Bender, D. B. (1972). Visual properties of neurons in inferotemporal cortex of the macaque. *Journal of Neurophysiology*, 35, 96–111.
- Haken, H. (1977). *Synergetics: An Introduction*. Springer.
- Hebb, D. O. (1949). *The Organization of Behavior*. John Wiley & Sons.
- Hirabayashi, T., Takeuchi, D., Tamura, K., & Miyashita, Y. (2013). Microcircuits for hierarchical elaboration of object coding across primate temporal areas. *Science*, 341, 191–195.
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
- Hopfield, J. J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the USA*, 79, 2554–2558.
- Hughes, J. R. (1995). The phenomenon of travelling waves: A review. *Clinical Electroencephalography*, 26(1), 1–6.
- Jansen, B. H., & Rit, V. G. (1995). Electroencephalogram and visual evoked potential generation in a mathematical model of coupled cortical columns. *Biological Cybernetics*, 73(4), 357–366.

- Kayser, C., Salazar, R. F., & König, P. (2003). Responses to natural scenes in cat V1. *Journal of Neurophysiology*, 90, 1910–1920.
- Köhler, W. (1930). *Gestalt Psychology*. Bells and Sons.
- König, P., & Schillen, T. B. (1991). Stimulus-dependent assembly formation of oscillatory responses: I Synchronization. *Neural Computation*, 3, 155–166.
- Kreiter, A. K., & Singer, W. (1996). Stimulus-dependent synchronization of neuronal responses in the visual cortex of the awake macaque monkey. *The Journal of Neuroscience*, 16(7), 2381–2396.
- Lachaux, J.-P., Rodriguez, E., Martinerie, J., & Varela, F. J. (1999). Measuring phase synchrony in brain signals. *Human Brain Mapping*, 8(4), 194–208.
- Lazar, A., Lewis, C., Fries, P., Singer, W., & Nikolic, D. (2021). Visual exposure enhances stimulus encoding and persistence in primary cortex. *Proceedings of the National Academy of Sciences of the USA*, 118(43), e2105276118, 1–11.
- Linnainmaa, S. (1970). *The representation of the cumulative rounding error of an algorithm as a Taylor expansion of the local rounding errors*. University.
- Livingstone, M. S. (1996). Oscillatory firing and interneuronal correlations in squirrel monkey striate cortex. *Journal of Neurophysiology*, 75(6), 2467–2485.
- Llinas, R. R. (1988). The intrinsic electrophysiological properties of mammalian neurons: Insights into central nervous system function. *Science*, 242, 1654–1664.
- Löwel, S., & Singer, W. (1992). Selection of intrinsic horizontal connections in the visual cortex by correlated neuronal activity. *Science*, 255, 209–212.
- Lowet, E., Roberts, M. J., Bonizzi, P., Karel, J., & De Weerd, P. (2016). Quantifying neural oscillatory synchronization: A comparison between spectral coherence and phase-locking value approaches. *PLOS ONE*, 11(1), e0146443, 1–37.
- Lowet, E., Roberts, M. J., Gips, B., De Weerd, P., & Peter, A. (2017). A quantitative theory of gamma synchronization in macaque V1. *elife*, 6, e26642, 1–44. (elifesciences.org)
- Lu, Z., Kim, J. Z., & Bassett, D. S. (2020). Supervised chaotic source separation by a tank of water. *Chaos*, 30(021101), 1–8.
- Lukoševičius, M., & Jaeger, H. (2009). Reservoir computing approaches to recurrent neural network training. *Computer Science Review*, 3(3), 127–149.
- Lundqvist, M., Rose, J., Herman, P., Brincat, S. L., Buschman, T. J., & Miller, E. K. (2016). Gamma and beta bursts underlie working memory. *Neuron*, 90(1), 152–164.
- Magee, J. C., & Johnstone, D. (1997). A synaptically controlled, associative signal for Hebbian plasticity in hippocampal neurons. *Science*, 275, 209–213.
- Maksymov, I. S., Huy Nguyen, B. Q., Pototsky, A., & Suslov, S. (2022). Acoustic, phononic, brillouin light scattering and faraday wave-based frequency combs: Physical foundations and applications. *Sensors*, 22(3921), 1–45.
- Markov, N. T., Ercsey-Ravasz, M., Van Essen, D. C., Knoblauch, K., Toroczkai, Z., & Kennedy, H. (2013). Cortical high-density counterstream architectures. *Science*, 342(6158), 1238406, 1–14.
- Markram, H., Lübke, J., Frotscher, M., & Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs. *Science*, 275, 213–215.
- McCulloch, W. S., & Pitts, W. (1943). A logical calculus of the ideas immanent in nervous activity. *The Bulletin of Mathematical Biophysics*, 5(4), 115–133.
- Milner, P. M. (1992a). The functional nature of neuronal oscillations. *Trends in Neurosciences*, 15, 387.
- Milner, P. M. (1992b). Reply. *Trends in Neurosciences*, 15, 387–388.
- Moy, W., Ahmed, I., Chiu, P.-W., Moy, J., Sapatnekar, S. S., & Kim, C. H. (2022). A 1,968-node coupled ring oscillator circuit for combinatorial optimization problem solving. *Nature Electronics*, 5(5), 310–317.
- Muller, L., Chavane, F., Reynolds, J., & Sejnowski, T. J. (2018). Cortical travelling waves: Mechanisms and computational principles. *Nature Reviews Neuroscience*, 19, 255–268.
- Murray, J. D., Bernacchia, A., Freedman, D. J., Romo, R., Wallis, J. D., Cai, X., Padoa-Schioppa, C., Pasternak, T., Seo, H., Lee, D., & Wang, X.-J. (2014). A hierarchy of intrinsic timescales across primate cortex. *Nature Neuroscience*, 17(12), 1661–1663.
- Nicolis, G., & Prigogine, I. (1977). *Self-Organization in Nonequilibrium Systems*. Wiley-Interscience.
- Nikolic, D., Häusler, S., Singer, W., & Maass, W. (2009). Distributed fading memory for stimulus properties in the primary visual cortex. *PLOS Biology*, 7(12), e1000260, 1–19.
- Offray de la Mettrie, J. (1748). *L'Homme Machine*. De l'Imp. d'Elie Luzac, fils. Leyde.
- Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381, 607–609.
- Onorato, I., Neuenschwander, S., Hoy, J., Lima, B., Rocha, K.-S., Broggin, A. C., Uran, C., Spyropoulos, G., Klon-Lipok, J., Womelsdorf, T., Fries, P., Niell, C., Singer, W., & Vinck, M. (2020). A distinct

- class of bursting neurons with strong gamma synchronization and stimulus selectivity in monkey V1. *Neuron*, 105(1), 180–197.
- Pascanu, R., Mikolov, T., & Bengio, Y. (2013). On the difficulty of training recurrent neural networks. *Proceedings of the 30th International Conference on Machine Learning, PMLR*, 28(3), 1310–1318.
- Peron, S., Pancholi, R., Voelcker, B., Wittenbach, J. D., Olafsdottir, H. F., Freeman, J., & Svoboda, K. (2020). Recurrent interactions in local cortical circuits. *Nature*, 579, 256–259.
- Quiroga, R. Q., Reddy, L., Kreiman, G., Koch, C., & Fried, I. (2005). Invariant visual representation by single neurons in the human brain. *Nature*, 435, 1102–1107.
- Rao, R. P., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2, 79–87.
- Ray, S., & Maunsell, J. H. R. (2015). Do gamma oscillations play a role in cerebral cortex? *Trends in Cognitive Sciences*, 19(2), 78–85.
- Ray, S., & Maunsell, J. H. R. (2010). Differences in gamma frequencies across visual cortex restrict their possible use in computation. *Neuron*, 67, 885–896.
- Rosenblatt, F. (1958). The perceptron. A probabilistic model for information storage and organization in the brain. *Psychological Reviews*, 65, 386–408.
- Salinas, E., & Sejnowski, T. J. (2001). Correlated neuronal activity and the flow of neural information. *Nature Reviews Neuroscience*, 2(8), 539–550.
- Schillen, T. B., & König, P. (1994). Binding by temporal structure in multiple feature domains of an oscillatory neuronal network. *Biological Cybernetics*, 70, 397–405.
- Schillen, T. B., & König, P. (1991). Stimulus-dependent assembly formation of oscillatory responses: II. Desynchronization. *Neural Computation*, 3, 167–178.
- Singer, W. (2021). Recurrent dynamics in the cerebral cortex: Integration of sensory evidence with stored knowledge. *Proceedings of the National Academy of Sciences of the USA*, 118(33), e2101043118, 1–12.
- Singer, W. (2019). A naturalistic approach to the hard problem of consciousness. *Frontiers in Systems Neuroscience*, 13(58), 1–9.
- Singer, W. (2018). Neuronal oscillations: Unavoidable and useful? *European Journal of Neuroscience*, 48(7), 2389–2398.
- Singer, W. (1999). Neuronal synchrony: A versatile code for the definition of relations? *Neuron*, 24, 49–65.
- Singer, W. (1993). Synchronization of cortical activity and its putative role in information processing and learning. *Annual Review of Physiology*, 55, 349–374.
- Singer, W., & Gray, C. M. (1995). Visual feature integration and the temporal correlation hypothesis. *Annual Review of Physiology*, 18, 555–586.
- Spyropoulos, G., Saponati, M., Dowdall, J. R., Schölvinc, M. L., Bosman, C. A., Lima, B., Peter, A., Onorato, I., Klon-Lipok, J., Roese, R., Neuenschwander, S., Fries, P., & Vinck, M. (2022). Spontaneous variability in gamma dynamics described by a damped harmonic oscillator driven by noise. *Nature Communications*, 13(2019), 1–18.
- Stryker, M. P. (1989). Is grandmother an oscillation? *Nature*, 338, 297–298.
- Swadlow, H. A. (1985). Physiological properties of individual cerebral axons studied *in vivo* for as long as one year. *Journal of Neurophysiology*, 54(5), 1346–1362.
- Tsao, D. Y., Freiwald, W. A., Tootell, R. B. H., & Livingstone, M. S. (2006). A cortical region consisting entirely of face-selective cells. *Science*, 311, 670–674.
- Van Rullen, R., Guyonneau, R., & Thorpe, S. J. (2005). Spike times make sense. *Trends in Neurosciences*, 28(1), 1–4.
- von der Malsburg, C., & Buhmann, J. (1992). Sensory segmentation with coupled neural oscillators. *Biological Cybernetics*, 67, 233–242.
- von der Malsburg, C., & Schneider, W. (1986). A neural cocktail-party processor. *Biological Cybernetics*, 54, 29–40.
- Wang, Y., Markram, H., Goodman, P. H., Berger, T. K., Ma, J., & Goldman-Rakic, P. S. (2006). Heterogeneity in the pyramidal network of the medial prefrontal cortex. *Nature Neuroscience*, 9(4), 534–542.
- Waxman, S. G. (1980). Determinants of conduction velocity in myelinated nerve fibers. *Muscle & Nerve*, 3(2), 141–150.
- Werbos, P. J. (1982). Applications of advances in nonlinear sensitivity analysis. Drenick, R. F., & Kosin, F., Editors. *System Modeling an Optimization*. Springer-Verlag, New York.