

RESEARCH ARTICLE

Neuronal networks quantified as vector fields

Szilvia Szeier, Henrik Jörntell ^{*}

Department of Experimental Medical Science, Lund University, Lund, Sweden

^{*} henrik.jorntell@med.lu.se

Abstract

The function of the brain is defined by the interactions between its neurons. But these neurons exist in tremendous numbers, are continuously active, and densely interconnected. Thereby, they form one of the most complex dynamical systems known and there is a lack of approaches to characterize the functional properties of such biological neuronal networks without resorting to dimensionality reduction methods. Here, we introduce an approach to describe these functional properties by using its activity-defining constituents, the weights of the synaptic connections and the current activity of its neurons. We show how a high-dimensional vector field, which describes how the activity distribution across the neuron population is impacted at each instant of time, naturally emerges from these constituents. We show why a mixture of excitatory and inhibitory neurons and a diversity of synaptic weights are critical to obtain a network vector field with a structural richness. We argue that this structural richness can be the foundation of achieving the diverse, dynamic activity patterns across the neuron population observed in recordings *in vivo* and thereby an underpinning of the behavioral flexibility and adaptability that characterizes biological creatures.

 OPEN ACCESS

Citation: Szeier S, Jörntell H (2025) Neuronal networks quantified as vector fields. PLOS Complex Syst 2(5): e0000047. <https://doi.org/10.1371/journal.pcsy.0000047>

Editor: Juan Gonzalo Barajas-Ramirez, IPICYT: Instituto Potosino de Investigacion Cientifica y Tecnologica AC, MEXICO

Received: September 23, 2024

Accepted: April 14, 2025

Published: May 19, 2025

Peer Review History: PLOS recognizes the benefits of transparency in the peer review process; therefore, we enable the publication of all of the content of peer review and author responses alongside final, published articles. The editorial history of this article is available here: <https://doi.org/10.1371/journal.pcsy.0000047>

Copyright: © 2025 Szeier, Jörntell. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data availability statement: All data are in the manuscript and/or supporting information files.

Funding: This work was funded by Vinnova Sweden (HJ) (grant# 2022-00943), Vetenskapsrådet (HJ) (grant#2019-01623,

Author summary

Understanding the brain with its densely interconnected network at its full complexity has been the subject of decades of research in neuroscience. Starting from studying the activity of just individual neurons, the field has migrated to investigating the collaboration of populations of neurons. Even though it is now understood that brain functionalities are a result of interactions between a high number of neurons, many approaches resort to describing neuronal interactions with a low number of properties. We argue that the high number of neurons require a high-dimensional representation in order to retain essential information about the population-level behavior. Here we introduce a neuronal network representation, namely in the form of vector fields, that is an inevitable consequence of how neurons interact and impact each other. These vector field representations, which generalize to arbitrary number of dimensions, dictate the temporal evolution of neuron population-level activity. We show what

grant#2023-03005) and Hjärnfonden (HJ) (grant# FO2024-0405-HK-126). The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests: The authors have declared that no competing interests exist.

network properties give rise to diverse and dynamic vector fields that govern neuron population-level activity reminiscent of experimental observations *in vivo*.

Introduction

Understanding the collaboration between multiple neurons remains a challenging issue within neuroscience. An important part of the difficulty is that the networks within the brain are often, at least to some extent, recurrently connected, such as in the neocortex [1–7] and in the spinal cord, both internally and externally via the sensorimotor feedback loops [8]. The specific function of any given neuron will then depend on the current activity in other neurons. Therefore, the operation of the brain networks can be nearly impossible to understand by observing activity neuron-by-neuron. Understanding brain function may then be improved by examining the mechanisms that control how the activity distributions across neuron populations evolve. Dynamical systems theory could offer tools for analyzing these mechanisms.

In biomedicine, dynamical system theory has previously been applied to the analysis of the propagation of activity in the heart [9,10] and macroscopically in the brain [11]. In both cases, the application has been to provide descriptive approximations of the observed propagations of the activity across the physical 2D or 3D space of the tissue in the form of vector fields. The vector field indicates, for example, that when the activity has reached a certain point in the heart tissue, it is more likely to propagate in the direction where the vector points and large deviations might indicate cardiac disease.

Rather than arbitrarily constructing vector fields from observed activity propagations, here we aimed instead to derive vector fields as an emergent property from the synaptic interactions between the neurons. This could help explain the mechanisms and factors that control the activity propagation across the neuron population. Instead of dealing with activity propagation across the 3D physical tissue space, we then need to consider the high-dimensional activity space of the neuron population. Each neuron's activity level forms a separate dimension of the activity space of the network, which consequentially has as many dimensions as there are neurons. The location of the combined neuron population activity within that space corresponds to the 'state' of the network. The activity space can therefore instead be referred to as the state space of the network. The vector field describes the tendency of the neuron population activity to be pushed in a certain direction within its high-dimensional state space depending on where in its state space it is currently located. A good analogy can be to think of a ferric particle in a magnetic field. The magnetic field, or the force vector field that it forms, may point in a certain direction towards which the ferric particle will be driven.

The location within the network state space and the evolution of spatiotemporal activity patterns across the neuron population may reflect the brain's current "thoughts", its expectations of internal and external information, or the adjustments needed to update its behavioral trajectory. Concretely, a behavioral choice ultimately determines which sequences of muscle activation patterns are engaged and when they are executed. The brain's evolving neuronal activity patterns can directly translate into the formation of sequences of spatial muscle activation patterns [8,12–14], while "thinking" may be defined as a process deciding when and how to generate specific activation patterns.

In multi-unit recordings *in vivo*, it is typically observed that the neuron population activity follows certain constraints (called 'neural manifolds' in [15]). Such constraints can also be observed in intracellular recordings of the sensory responses in individual cortical neurons,

which tend to fall into specific categories as a result of the dynamic interaction between sensory input and the ongoing state evolution of the cortical network [5,6] in a manner compatible with the neural manifold idea. Different approximative descriptions of those observed constraints have been published to account for the propagation of the population activity (reviewed by [15,16]). Here, our objective is instead to describe the mechanisms by which such constraints on the travel across the network activity state space can arise. The phenomenon implies that there is some ‘hidden force’ that makes it less likely for the population to reach certain activity combinations [17]. We believe that the vector field that emerges from the synaptic interactions across the neuron population, as we describe in the present paper, could be an important factor of this ‘force field’, i.e. the mechanism that creates these constraints.

Current state-of-the-art approaches to extract information from multi-neuron recordings typically rely on performing some type of dimensionality reduction on the neuron population activity data [16,18]. But dimensionality reduction could result in a loss of the underlying information present in the brain circuitry signals, which could be critical to understand the neuronal interactions of the brain circuitry processing [19,20]. The potential problem with dimensionality reduction is the risk of greatly underestimating the complexity of the network operation. High complexity is likely necessary to allow the animal to juggle multiple functions and processes at the same time. For instance, the animal may need to blink, lick, make eye movements [21] to monitor unpredicted events, adjust body position and muscle tone distribution, and continuously plan ahead for the optimal timing of these and other adjustments. In addition, it may prepare for future behavioral decisions in response to changing conditions [19].

There is another potential problem with many current approaches to analyze neuron population activity. Jazayeri and Ostojic [18] compared different dimensionality reduction methods, which could report different results from the same underlying data. Put in other words, if we don’t know the computational architecture of the brain circuitry, it is not possible to get a dimensionality reduction right, or ‘loss-less’, regardless of the embedding method or dimensionality reduction tool used. Hence, there is a need for a conceptual framework to quantify neuron population level interactions, while not discarding essential parts of the information due to arbitrary impacts of the particular frame of reference (‘embedding’) used to approximate the nature of those interactions.

Our focus here is to identify a conceptual framework that would be an emerging property of the fact that synaptically connected neurons will inevitably impact each other’s activity levels and that can serve as a basis for new tools to characterize and understand multi-neuron interactions. We introduce the method of describing the activity- and weight-dependent neuronal interactions using vector field representations. We show how this representation can be used to understand the factors that impact the structure of the state space of the neuronal network, regardless of its dimensionality (number of neurons). We show how the distribution of the synaptic weights across the network shapes its vector field and that inhibitory neurons are important to achieve a diversity of operational states within the network state space. We also demonstrate how subspaces, specific planes in the high-dimensional vector field, can be extracted for visualization purposes and to explore in greater detail factors that would impact the multi-neuronal interactions in specific contexts, such as in a specific movement phase or in more abstract representations of behavioral choices.

Materials and methods

Non-spiking neuron model

To build networks in which the underlying principles of control of the activity distributions across a neuron population could be examined, we utilized a previously published non-spiking neuron model called the Linear Summation model (LSM) [22]. This model emulates a conductance-based neuron, which is derived from the Hodgkin-Huxley (H-H) neuron model [23]. The activity of the post-synaptic neuron depends on the activities of the presynaptic neuron(s) and can be formulated as

$$a_j = \frac{\sum(w_i * a_i)}{k + \sum|w_i * a_i|}, \quad (1)$$

where a_j is the activity of the postsynaptic neuron, a_i represents the activity of a presynaptic neuron, w_i is its synaptic weight, and k is the static leak.

Notably, this implies that the variable a_i is interpreted as the time-averaged firing rate of the neuron i . Changes in the firing rate of the presynaptic neurons are then the main source of change that can impact the firing rate of the postsynaptic neuron. The modifications needed to adapt the vector field representation for spiking neural networks and other subcellular nonlinearities are discussed later in [Results](#) and [Discussion](#). Here, however, we focus on introducing the fundamental elements of the framework and outlining its core principles.

In a biological neuron, due to its electrical signaling being driven by opening and closing conductances (ion channels), the effect of a given synapse's activation depends not only on its synaptic weight, but also on the number of other conductances open at the same time. The LSM represents two types of such conductances: those of other synapses active at the same time and the static leak conductance. The latter represents the constitutive leak ion channels that establish and maintain the resting membrane potential and that therefore in the biological neuron always need to be open. The output activity of the LSM neuron is calculated from the weighted sum of the presynaptic neural activity, as shown in Eq 1. The numerator accounts for the linear combination of the incoming activity from presynaptic connections scaled by their respective synaptic weights. A natural consequence of conductance-based signaling in the biological neuron is that the higher the activity of all other synapses (or conductances) on the neuron, the lower the impact of a given synapse with a given weight (i.e., a given conductance). As a result, in a larger neuron, a synapse of a given weight will have less impact on the postsynaptic neuron compared to the same synapse on a smaller neuron. Notably, this also contributes towards a normalization, or a prevention of saturation, of the output activity, which in the LSM is captured by the denominator, where the $\sum|w_i * a_i|$ term is the total magnitude of synaptic input and k denotes the static leak of the neuron. Since the number of constitutive leak channels within a unit dendritic area is approximately constant [24,25] and the membrane area scales with the number of synapses the neuron receives, k is here calculated as a constant multiplied by the number of synapses on the neuron. For further detailed explanation and derivation of the model, readers are referred to the original work [22]. In our experiments, we modeled fully connected networks so the number of synaptic connections on a given neuron was always one less than the number of neurons. Therefore, unless stated otherwise, k was $0.1 \times (n - 1)$, where n is the number of neurons in the network.

Neural state space

The impact exerted by one (presynaptic) neuron on another (postsynaptic) neuron is proportional to the presynaptic activity and synapse weight. Synaptic weights were set to $[0 \dots 1]$ for excitatory synapses and to $[-1 \dots 0]$ for inhibitory neurons; autapses and parallel synapses were not allowed. A weight of 0 can be interpreted as a silent synapse or the absence of synaptic connectivity. The activity of individual neurons is bounded between 0 (no activity) and 1 (saturation or epilepsy). We can define a (bounded) state space where each dimension represents an individual neuron's activity level. Such a state space can be constructed for a network of arbitrary dimensionality. Essentially, we obtain a hypercube with the dimensionality being equal to the number of neurons in the network. Every state (point) in the state space, represents a neural activity distribution that uniquely determines the activity configuration of all the neurons in the network.

Two-dimensional planar subspace

We can consider two-dimensional, planar subspaces (or extracted planes) of the original, high-dimensional neural state space. The two dimensions of these subspaces, or axes of the plane, represent the activity of neurons: either two individual neurons at a time or a population of neurons divided into two groups, one per dimension. When the axes represent individual neurons, the extracted plane reflects the relationship between those two neurons. However, an axis of the extracted plane can also represent a linear combination of multiple neurons. In this case, the extracted plane illustrates the relationship between the two chosen linear combinations of neurons within the network. A neuron is considered to be represented on the extracted plane if its dimension in the neural state space is not orthogonal to the plane, referred to as a “non-perpendicular” neuron. In contrast, neurons whose dimensions are orthogonal to the chosen plane are called “perpendicular” neurons.

Extracting a plane from the neural state space enables the analysis of network components, such as the influence of a specific neuron's activity level on the neuron population represented along the axes of the extracted plane. To achieve this, the plane should be selected such that the neuron, or neuron population, whose impact is being studied has dimensions in the neural state space that are orthogonal to the plane. Meanwhile, the neuron population affected by these interactions should be represented along the axes of the extracted plane.

It is indeed possible to select a plane such that no neuron is perpendicular. However, if the studied population was represented on the axes of the plane, the vector field in that plane would reflect the changes applied to the studied population instead of the effect of those changes on the target population.

Critical point location (Derivation of Eq 9)

In our vector field representation, critical points are activity level combinations where there is a match between the weighted presynaptic excitation and inhibition. Given networks with uniform weights, the critical point is always located along the central diagonal. In order to displace the critical point from the central diagonal, the synaptic weight distribution must be skewed.

The critical points are located where vector lengths are zero, or alternatively, $\forall j \in [1 \dots n] : a_j = 0$, where n is the number of neurons and a_j are the vector components obtained from Eq 1. When equating Eq 1 to zero, we can simplify the formula to

$$a_j' = \sum (w_i * a_i). \quad (2)$$

Given a 3-neuron network, if we extract a plane that is perpendicular to N3 at activity level a_3 , the location of the critical point is determined by the weights of both perpendicular and non-perpendicular neurons. A general rule for the location of the critical point within the extracted plane, in terms of activity level, can be calculated as

$$\begin{cases} a_1' = w_3 a_2 + w_1 a_3 = 0 \\ a_2' = w_4 a_1 + w_2 a_3 = 0 \end{cases}, \quad (3)$$

where a_i are the activity levels of the neurons and w_1, w_2 are the synaptic weights connecting N3 to N1 and N2, respectively. The synapse from N2 to N1 has weight w_3 , while the reverse connection has weight w_4 . We can rearrange the equations as

$$\begin{cases} w_3 a_2 = -w_1 a_3 \\ w_4 a_1 = -w_2 a_3 \end{cases}. \quad (4)$$

Next, we divide the top equation with the bottom one and simplify.

$$\frac{w_3 a_2}{w_4 a_1} = \frac{-w_1 a_3}{-w_2 a_3} = \frac{w_1}{w_2}. \quad (5)$$

Finally, we can express the critical point location in terms of an activity level ratio of N1 and N2 (the non-perpendicular neurons) as

$$\frac{a_2}{a_1} = \frac{w_1 w_4}{w_2 w_3}. \quad (6)$$

A ratio of 1 corresponds to a critical point that is located along the central diagonal. Therefore, it is evident from Eq 6 that a skewed synaptic weight distribution is necessary to displace the critical point from the central diagonal.

Results

Vector field representation and plane extraction

We started out with a setting of synaptically connected excitatory cortical neurons (Fig 1A) to introduce the concept of how neuronal interactions can be quantified as vector fields. As indicated in Fig 1A we used a non-spiking neuron model with identical, linear-like neuronal input-output functions (emulating the conductance-level biological neuron behavior [22]) across all neurons. Fig 1B illustrates a network with only two reciprocally connected excitatory neurons. If a neuron is connected to another neuron with a synapse, then that means that the activity of the second neuron is dependent on the activity of the first neuron, as shown in Fig 1C at $y = 0$ along the x-axis. Hence, the first neuron will increase the activity of the second neuron, if the synapse is excitatory. That impact will depend on the activity level of the first neuron - if that activity is very low or zero, then it will have no or very low impact on the second neuron, and vice versa for high activity. It should be noted that with the neuron model we employed, the series of vectors along the x-axis at $y=0$ increased at a decelerating rate since the higher the activity of all other synapses (or conductances) on the neuron, the lower the impact of additional inputs [22] (Fig 1C; see Methods).

At the same time, the activity of the second neuron will also impact the activity of the first neuron in a similar manner, as they are reciprocally connected. The interaction between

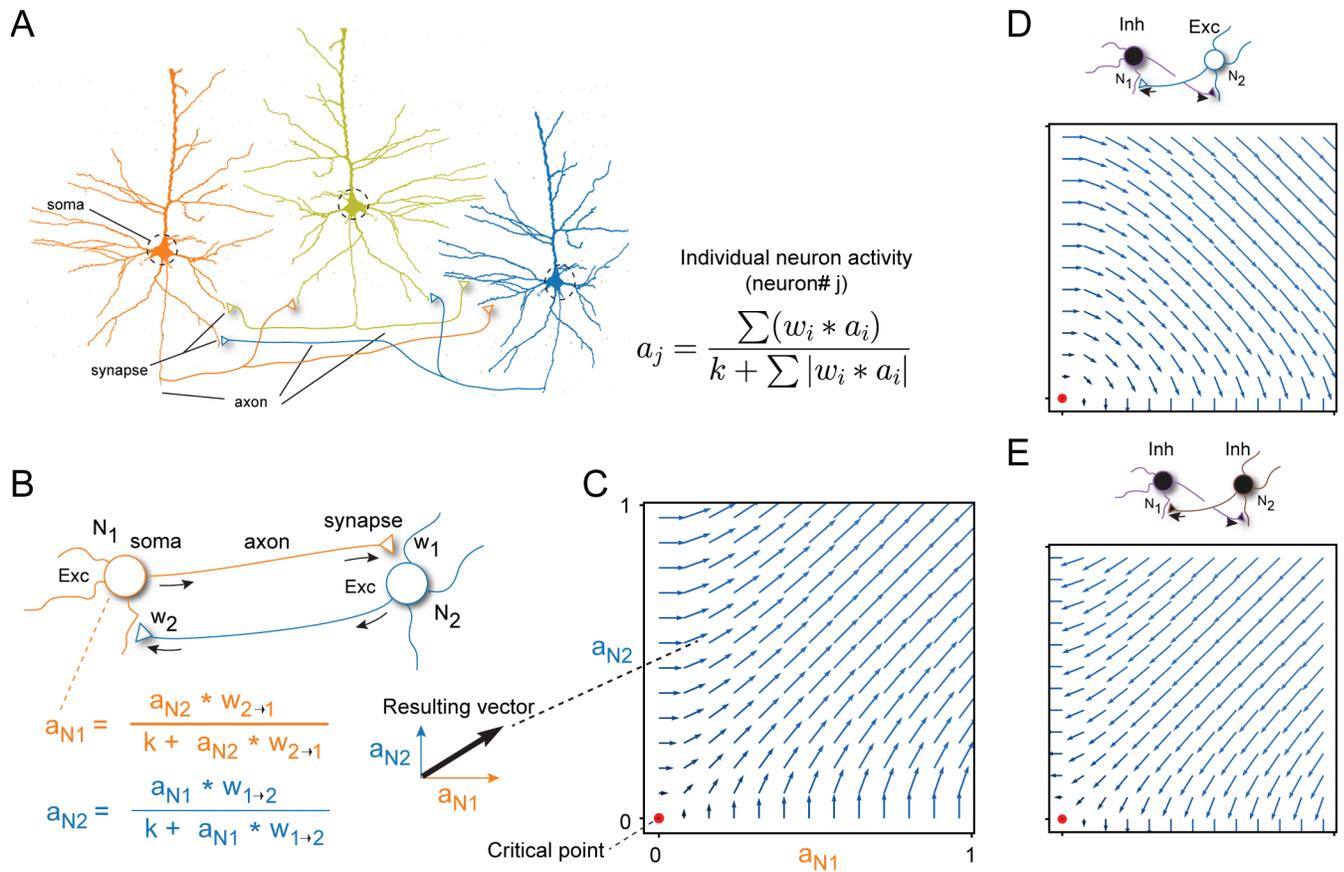


Fig 1. Mapping neuronal networks to vector field representations. (A) Actual neurons have dendrites, a soma, and an axon. The axon branches profusely and makes synapses on other neurons. The example shows outlines of three (excitatory) pyramidal neurons (dendrites and axons are truncated) from previous histological analysis, as well as putative synaptic connections between the three. The inset shows the neuron model used, i.e. how the activity of each neuron was calculated. (B) Representation of the network formed between two of the pyramidal neurons. Arrows indicate the flow of signals. At each activity level, we calculated a vector from the activity levels of the two neurons as indicated. (C) Resulting vector field representation of the the two-neuron network. Each axis represents the activity of one of the neurons. (D) Vector field representations of an excitatory-inhibitory network. (E) Same as D but for an inhibitory-inhibitory network. Across all vector fields, synaptic weights were uniformly set to 1 (negative sign for connections originating from inhibitory neurons) and *k* was set to 0.2.

<https://doi.org/10.1371/journal.pcsy.0000047.g001>

the two neurons creates a vector component indicating the magnitude that the two neurons impact each other, at any possible activity level (Fig 1B). This process can then be iterated for all activity combinations across the two neurons, each combination resulting in a vector component. The resulting vector field across all possible activity combinations for the two neurons is shown in Fig 1C. As the two neurons form a positive feedback loop, they will tend to push each other's activities up towards the upper right corner, i.e. this would correspond to saturation or overexcitation of the neuronal activity [26] (corresponding to a maximum spiking probability of 1, as illustrated for example in [27]; in biology this would correspond to an epileptic state). In cases where one or both neurons are instead inhibitory (Fig 1D, 1E), the structure of the vector field will reorganize accordingly. When both neurons are inhibitory (Fig 1E), the vectors point downwards relative to the two axes and would hence tend to drive the activity of both neurons towards zero. Notably, the only neutral position within any of these vector fields, i.e. the position where the vector lengths approach zero, is when both neurons have zero activity ('critical point') (Fig 1C–1E). In dynamical system theory, a

critical point indicates a subspace where the vectors have zero length without necessarily implying that the surrounding vectors point towards that point. Other related concepts are attractors and limit cycles, which are descriptive, approximative structures in the vector field that the vectors are pointing towards or away from. Notably, below we will iteratively refer to the critical points but merely as a proxy for changes in the overall vector field structure, because illustrating the whole set of vector fields would have resulted in less comprehensive figures. But we do not otherwise put great emphasis on critical points here.

When the number of neurons connected to each other increases, that means that the dimensionality of the network increases. First, consider a network of three excitatory neurons (Fig 2A). Each neuron will be impacted by the synaptic input activities of the other two neurons. Using the same approach as above (Fig 1C), we can calculate the now 3-dimensional vector field (Fig 2B). However, it is harder to make a good visualization of the vector field in 3D, and if we add more neurons to the network, it becomes impossible. For visualization purposes, we can extract a two-dimensional subspace (i.e. a plane) from the full 3-dimensional network space to visualize the interactions between two selected neurons (Fig 2C), which will now look analogous to the 2-dimensional networks of Fig 1. Similarly to the 3-dimensional case, we can extract planes from a network consisting of 4 neurons (Fig 2D) and from 8-dimensional networks (Fig 2E). This concept of plane extraction can be generalized to networks of arbitrary dimension.

Choosing a plane for visualization is equivalent to fixing the activity of all neurons orthogonal to the axes of the selected plane. We refer to the set of activity-fixed neurons as the

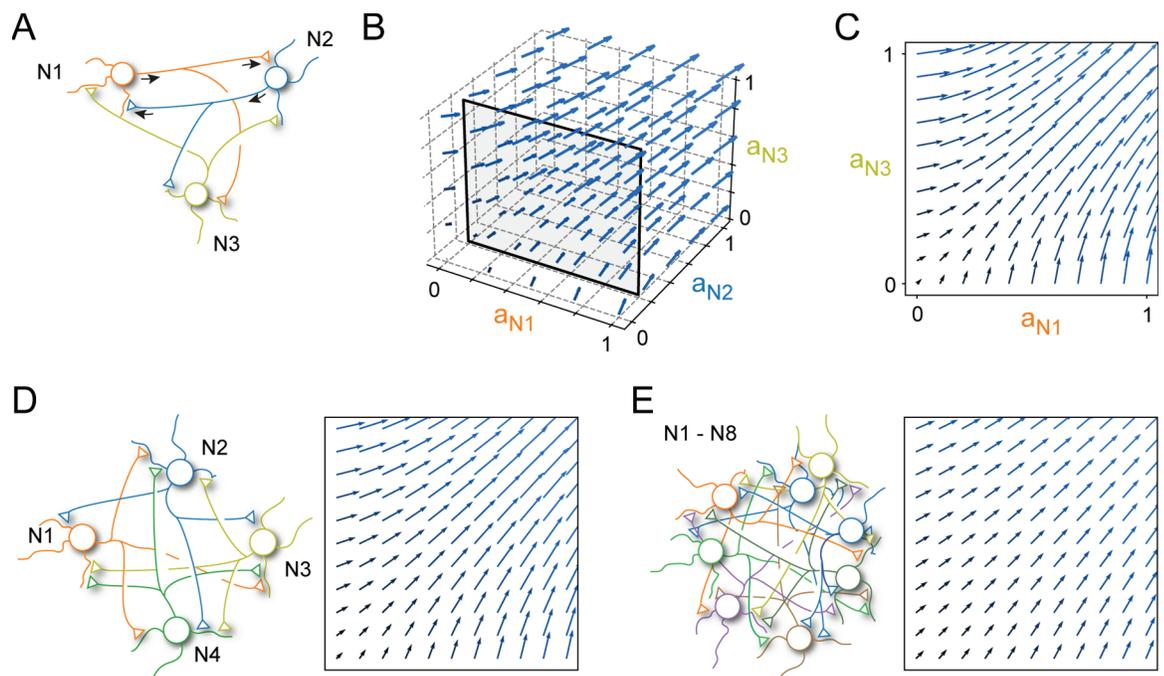


Fig 2. Extending the vector field representation to higher dimensions. (A) Representation of the excitatory network formed between the three pyramidal neurons. (B) Vector field representation of the three-neuron network. (C) A 2D plane extracted from the 3D vector field, corresponding to the plane where the activity of the perpendicular neuron (N2) was fixed. (D) Network consisting of four pyramidal neurons and an example plane extracted vector field from the 4-dimensional vector field space. (E) An eight pyramidal neuron network with an example extracted plane. Across all panels, the synaptic weights were uniformly set to 0.1. The perpendicular neurons were set to have an activity of 0.1.

<https://doi.org/10.1371/journal.pcsy.0000047.g002>

‘perpendicular neurons’ since their activity axes are perpendicular to the visualized plane. If we vary the activity level of the perpendicular neuron(s), it will lead to different vector fields in the visualized plane, as the influence of the perpendicular neuron(s) on the other neurons will change with their activity level(s). Note that the critical point could end up being located outside the visualized plane, if any of the perpendicular neurons has non-zero activity, which indeed was the case for all planes illustrated in Fig 2.

The two dimensions of the visualized plane can also represent a linear combination of neurons in the network, rather than being limited to a single neuron per axis. In the same 3-neuron excitatory network as before (Fig 3A), we can orient the plane in its 3-dimensional vector field such that one of its axes represents the combined activities of two out of the three neurons, i.e. the axis is orthogonal to only one of the neurons (N3 in Fig 3B). The x-axis of the extracted plane now represents a weighted linear combination of the first two neurons (Fig 3C), with each neuron’s contribution determined by the angle of the plane’s orientation relative to the neuron activity axes. Since this axis now corresponds to the combined activity of multiple neurons, its maximum value can exceed 1, even though the activity of any single neuron cannot (Fig 3C, red arrows). This concept of plane extraction can be extended to

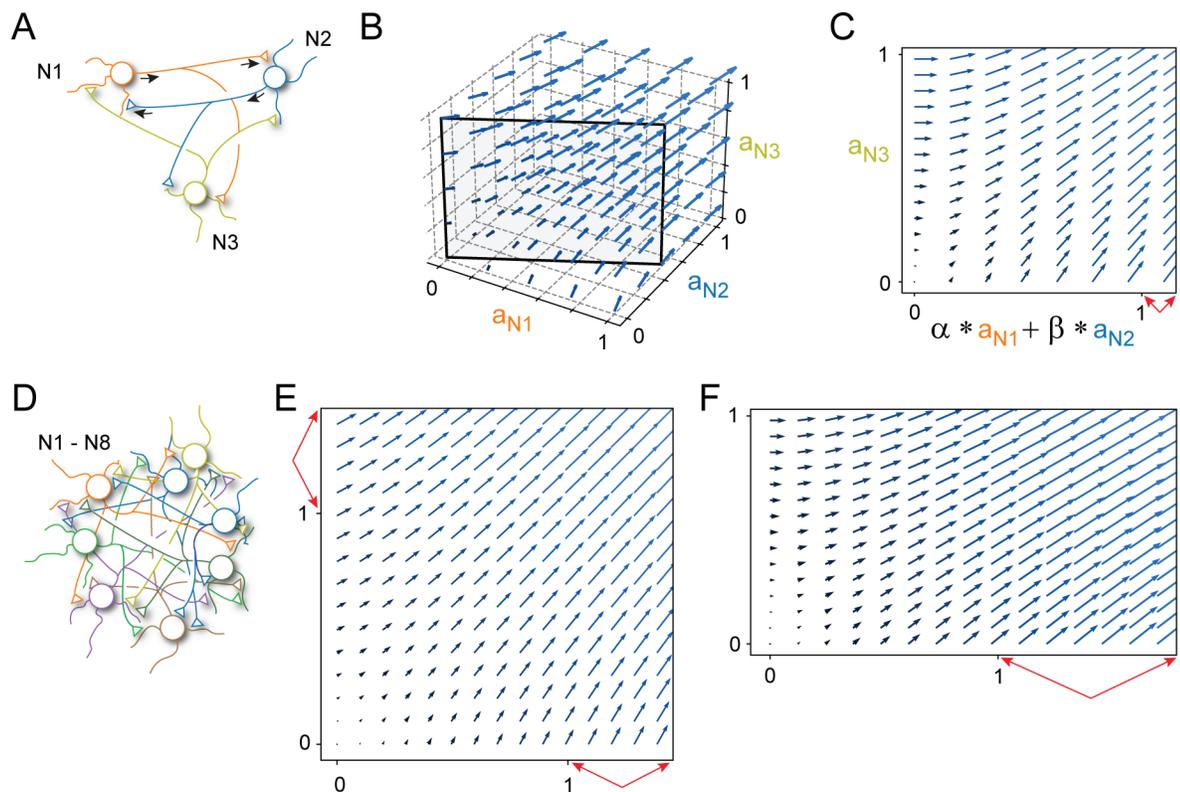


Fig 3. The axes of extracted planes can represent the activity of more than one neuron at a time. (A) The 3-dimensional excitatory network. (B) Vector field representation of the entire 3-dimensional state space with a plane which holds information about the activity of two out of the three neurons in one of its dimensions in its x-axis. (C) The 2D plane extracted from the vector field. The x-axis is a linear combination of two out of the three neurons as indicated ($\alpha = \frac{2}{\sqrt{5}}, \beta = \frac{1}{\sqrt{5}}$). The red arrows indicate that the x-axis, as a consequence, has a maximal value that exceeds 1. (D) An eight neuron excitatory network. (E) Extracted plane from the 8-dimensional network with two neurons represented on each axis (N1,N2 and N3,N4). (F) Plane extracted from the same network but where one axis represents the combined activity of 3 neurons while the other just a single one (N1-N3 and N4). Across all panels, the synaptic weights were uniformly set to 0.1, all perpendicular neurons had an activity of 0.

<https://doi.org/10.1371/journal.pcsy.0000047.g003>

networks of arbitrary dimension (Fig 3D), where each of the two dimensions (axes) of the plane can represent any combination of neurons (Fig 3E, 3F). In fact, a single neuron could even be represented on both dimensions of the extracted plane, if both axes of the plane are oriented at non-perpendicular angles to the activity axis of that neuron. Below, we will use this plane visualization to examine in detail how changes in network parameters, such as a specific synaptic weight, affect the vector field.

Location of the critical point depends on synaptic weights

A critical point is a point in the state space where no change in activity occurs (Fig 1). In visual terms, this means that the vector lengths at that position are zero. Due to the neuron model definition (Eq 1), any change in the synaptic weights or in the activity of the perpendicular neuron(s) can cause the critical point to shift locations, and this will be accompanied by structural changes in the surrounding vector field. We can analytically determine the location of the critical point by only taking the numerator of the neuron model (Eq 1) into consideration:

$$\hat{a}_j = \sum (w_i * a_i) = 0. \quad (7)$$

Given a network which is either fully excitatory or fully inhibitory, we can see that the only solutions will be

$$a_i = \begin{cases} [0 \dots 1], & \text{if } w_i = 0 \\ 0, & \text{otherwise} \end{cases}. \quad (8)$$

If the synaptic weights are 0, we are left with disconnected neurons. Therefore, in a fully excitatory or in a fully inhibitory network where the neurons are fully connected, the location of the critical point will be at zero (the origin). This motivated us to explore whether it would be possible to displace the critical point from the origin, if the neural network contains both excitatory and inhibitory neurons, as is typically observed in the brain.

First we consider a 3-neuron network consisting of 2 excitatory neurons and 1 inhibitory neuron (Fig 4A). If we make the inhibitory neuron perpendicular, then the vector field plane describes the interaction between N1 and N2, at a given selected activity level of the inhibitory neuron (N3). Hence, we made the inhibitory neuron the perpendicular neuron to study the impact it had on the vector field formed by the two excitatory neurons. The location of the critical point can now be offset from zero activity in the two neurons N1 and N2 (Fig 4B), which is due to the match of excitation and inhibition in those points. In this case, when we have outgoing synaptic connections of equal weight ($w_1=w_2$) from the inhibitory neuron to the two excitatory neurons, they will be impacted equally. Now, if we increase the weights of the inhibitory neuron (while maintaining the condition $w_1=w_2$), the location of the critical point will move only along the central diagonal (Fig 4B) in the neuron activity space of N1 and N2. The surrounding vector field also dramatically changes its structure, as shown by the two examples in Fig 4C.

We can also reverse the configuration, with N1 and N2 as inhibitory neurons and the perpendicular neuron (N3) as excitatory, while maintaining equal weights ($w_1 = w_2$) on N1 and N2 (Fig 4D–4F). Increasing the excitatory weights (while maintaining $w_1 = w_2$) has the same effect on the critical point as earlier, i.e. its location moves upwards along the central diagonal in the vector field (Fig 4B, 4E). Importantly, however, the structures of the vector fields, again provided as two examples, are 'inverted' for the inhibitory network (see the directions of the vectors; Fig 4C, 4F).

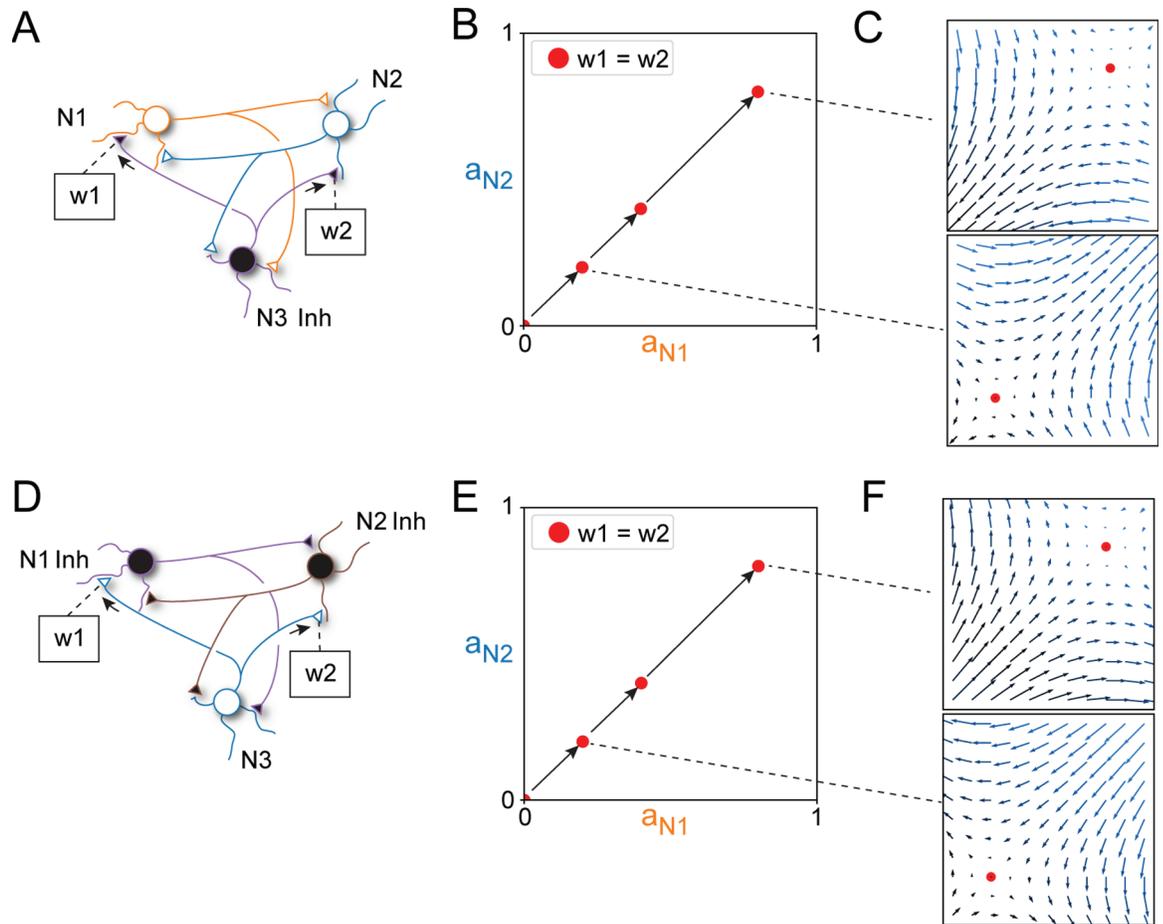


Fig 4. In networks with both excitatory and inhibitory neurons the critical point can be controlled. (A) A three-neuron network with one inhibitory neuron (Inh). The weights of the synapses made by the inhibitory neuron are indicated as w_1 and w_2 . (B) Extracted plane with the activity of the two excitatory neurons represented on the axes, and the inhibitory neuron represented on the perpendicular axis. When the weights w_1 and w_2 were increased from zero (both weights were always equal), the location of the critical point moved upwards along the central diagonal of the vector field (individual vectors are not shown for clarity). For the three illustrated points, the inhibitory weights were -0.2 , -0.4 and -0.8 . (C) Two examples of the full vector fields at different weight magnitudes of w_1 and w_2 . (D) A network of two inhibitory neurons and one excitatory neuron, where w_1 and w_2 were instead the weights of the excitatory synapses. (E) Similar effects as in B arise when the excitatory weights of the perpendicular neuron are increased. The synaptic weight magnitudes of the perpendicular neuron are the same as in B. (F) Examples of full vector fields. Across all panels the weights between N1 and N2 and the activity of the perpendicular neuron were fixed at 1.

<https://doi.org/10.1371/journal.pcsy.0000047.g004>

If the synaptic weights of the perpendicular neuron are not uniform, the critical point location will be displaced from the central diagonal because of the greater impact the perpendicular neuron has on one of the neurons. The magnitude of the displacement from the diagonal depends on the relative size of the outgoing synaptic weights: the larger the difference between those weights (the skewness), the larger the displacement from the central diagonal (Fig 5A, 5B). As a general rule, the skewness of the outgoing synaptic weights of the perpendicular neuron on the neurons represented on the respective axes is what determines the extent of displacement from the central diagonal.

However, the interactions are also influenced by the weights between the non-perpendicular neurons (the neurons represented on the axes). Previously, these weights were equal ($w_3 = w_4$) (Fig 5C). If we instead skew these weights, the movement of the critical point will

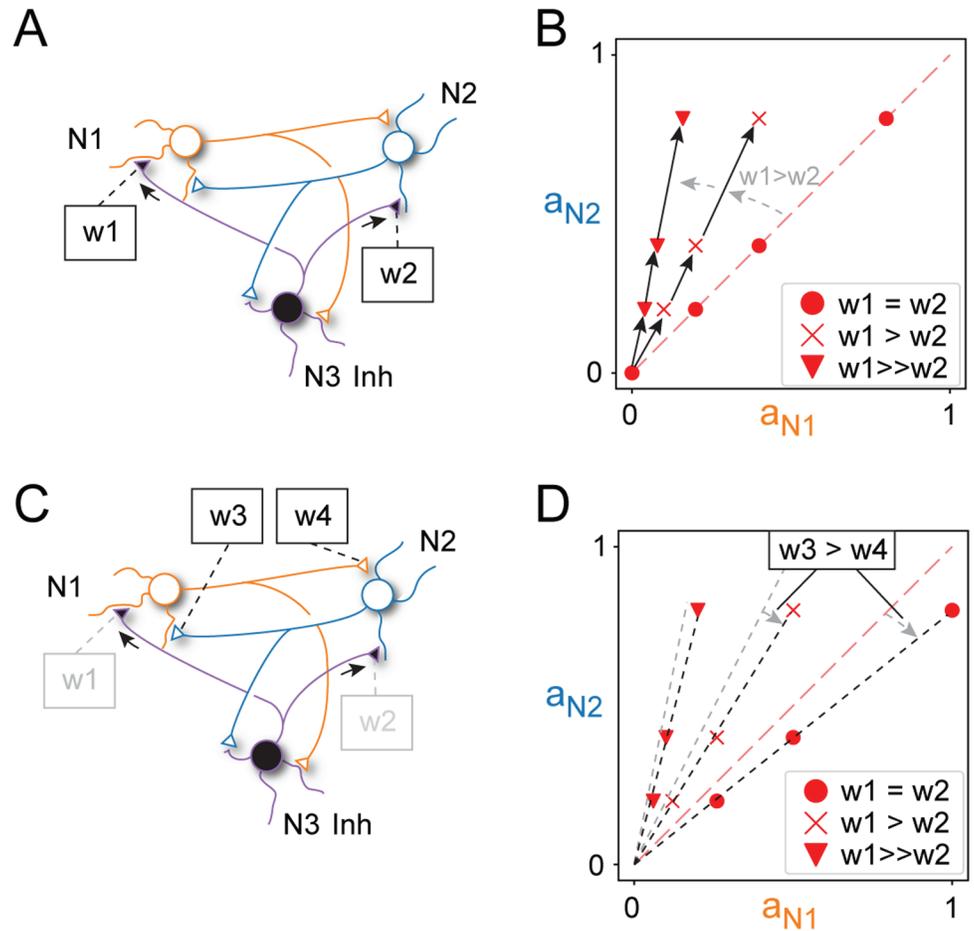


Fig 5. The location of the critical point can shift dramatically when the weights are non-uniform. (A) A network with two excitatory and one inhibitory neuron. Synaptic connections between N1 and N2 were uniformly set to 1. (B) The location of the critical point is offset from the central diagonal when the synaptic weights w_1 and w_2 , formed by the perpendicular neuron, are skewed rather than equal. Each symbol indicates a skewness level, a series of a symbol indicates the location of the critical point when the weights of w_1 and w_2 were increased proportionally. The weight magnitudes of w_1 were set to 0.2, 0.4, 0.8; w_2 was scaled down with the following proportions: 1 (unchanged), 2, 5. (C) The same network indicating the synaptic weights w_3 and w_4 for the synapses made between the non-perpendicular neurons. (D) Impact of making the weights w_3 and w_4 non-uniform ($w_3=1$, $w_4=0.8$). The arrows indicate the impact on the locations of the critical points when the weights were made non-uniform for the same skewness levels of w_1 and w_2 as in B. Across all panels the activity of the perpendicular neuron were fixed at 1.

<https://doi.org/10.1371/journal.pcsy.0000047.g005>

deviate from its original diagonal path observed when the weights w_1 and w_2 were proportionally increased (Fig 5D). Thus, the trajectory of the critical point, when varying the weights of the perpendicular neuron (w_1 , w_2), is also affected by the weight distribution between the two non-perpendicular neurons. In our 3-dimensional network example, the general relationship between the critical point location, in terms of neuron activity levels, and incoming synaptic weights can be expressed as

$$\frac{a_2}{a_1} = \frac{w_1 w_4}{w_2 w_3}, \tag{9}$$

where a_1 , a_2 are the activity levels of N1 and N2, the synaptic weights made by the perpendicular neuron N3 are w_1 , w_2 , while w_3 , w_4 are the weights between N1 and N2 (as shown in Fig 5C). For a derivation, see [Critical point location \(Derivation of Eq 9\)](#).

The vector field structure is directly linked to the structure of the neuronal network, which includes the synaptic weights as a central element (Fig 1). From Fig 5 we can conclude that when the synaptic weights dynamically change, such as in synaptic plasticity, the structure of the vector field adapts accordingly.

Location of the critical point can be controlled by the activity level

While synaptic weights shape the overall structure of the vector field, the specific activity levels of the neurons determine the subspace in which the network operates. Through manipulation of the perpendicular neuron activities, we can modify the position of the critical point and thus the structure of the underlying vector field within the extracted plane. The vector field structure is what impacts the evolution of the population level activity (i.e. the network activity state) in that plane, within the constraints that the synaptic weight distributions would allow for. This is illustrated in Fig 6, where we used a 4-neuron network with two perpendicular inhibitory neurons. The activities of the two perpendicular neurons were altered independently of each other to obtain different combinations of activity (Fig 6B). As shown in Fig 6C, the location of the critical point now travelled over a very large range of the activity space of the two non-perpendicular neurons, and the structure of the surrounding vector field also changed dramatically (Fig 6D).

Dependence of the vector field structure on the neuronal activation function

So far, we studied the structure of the vector field under the assumption of the specific input-output function of the LSM neuron model. The neuronal input-output function, sometimes also referred to as the f-I curve or the activation function, describes how much output (spike firing frequency) the neuron will generate in response to a range of summated synaptic inputs. The LSM builds on observations *in vivo*, where many neurons have a relatively linear activation function in their lower range of activity [28] and then become saturated at higher levels of activity due to the conductance-dependent shunting effect that we also call ‘protection leak’ in the LSM [22]. But other nonlinear activation functions have been observed in some neurons, and here we wanted to explore their possible impacts on the vector field of the network. Fig 7A first illustrates the decelerating activation function implemented by the LSM. As can be seen, the response function defines the vector component of that neuron by which it will impact its target neurons. Fig 7B illustrates this for an exponentially increasing, accelerating activation function, the output vectors of this neuron simply scale proportionally to that activation function. Some neurons may have an activation threshold, i.e. their baseline synaptic input is not sufficient to hold the neuron above its firing threshold and therefore a given synaptic input is not guaranteed to impact the neuronal firing rate (Fig 7C). The activation threshold simply shifts the range of input that will produce an output vector, and the vector component is again proportional to the output level. In this case, there will be a range of inputs for which the neuron does not provide any output (Fig 7D). Hence, whatever a neuron’s activation function, it is straightforward to include the impact of that activation function on the structure of the vector field. Even when the activation function dynamically changes, i.e. adaptive activation threshold [29–31], each change directly alters the structure of the vector field.

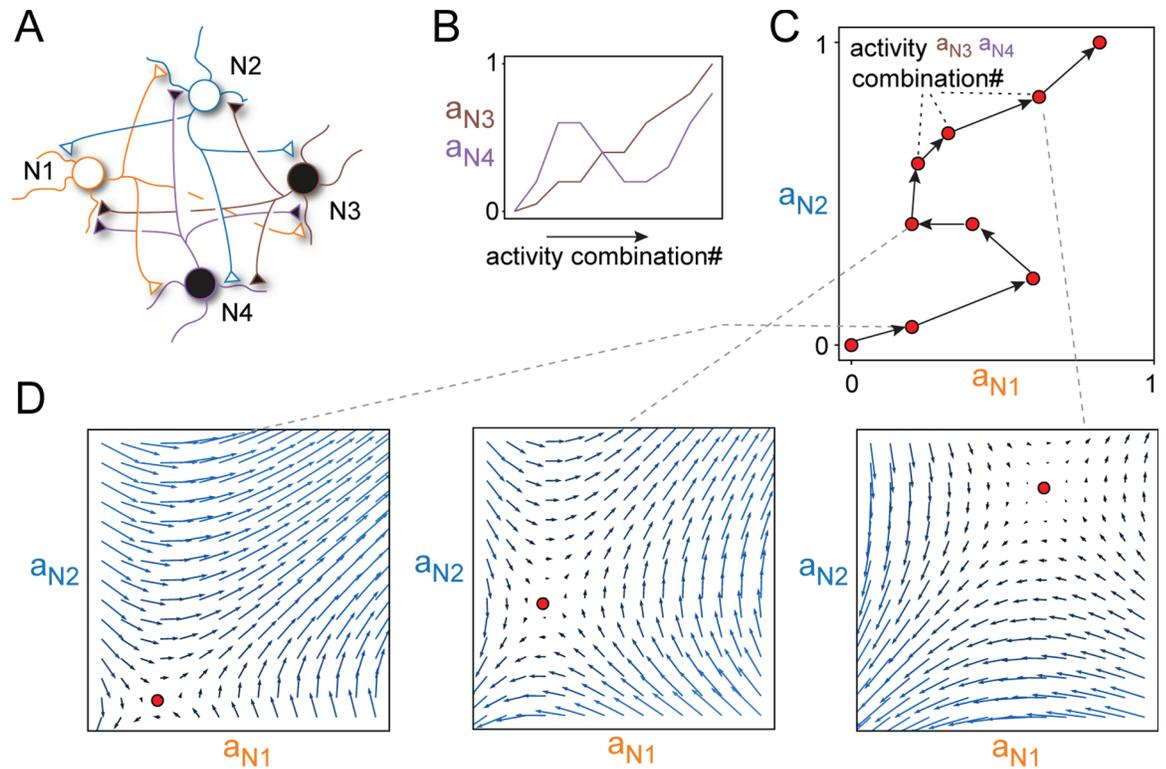


Fig 6. The activity levels of the perpendicular neurons dynamically control the critical point location in the extracted plane. (A) Four-dimensional network consisting of two excitatory and two inhibitory neurons. All synaptic weights were set to 0.5, except the synaptic weights from N3 to N2 and N4 to N1 which were reduced to 0.01. (B) Predefined activity level trajectories of the two inhibitory neurons being perpendicular to the extracted plane. (C) Critical point trajectory within the selected plane resulting from the activity level settings in B. (D) Examples of full vector fields from three of the locations of the critical point.

<https://doi.org/10.1371/journal.pcsy.0000047.g006>

Further, in cases where some neurons are only transiently activated, i.e. their output is silent for long periods of time and only occasionally do they erupt in an output, then the neuron will have no impact on the vector field structure while it is silent. This is similar to the situations in Figs 5 and 6, when the activity of the perpendicular neuron(s) is zero. The moment it becomes active, that neuron will exert its impact on the vector field depending on the intensity of its output activity (also shown in Figs 5 and 6).

Discussion

By introducing a conceptual framework to analyze neuron population level interactions within recurrent networks, we showed that vector fields inevitably emerge as a result of synaptic weights and neuronal input-output functions. The vector field will define the direction in which the network's overall activity state will evolve. In other words, the vector at any given point in the network state space will impact how the activity distribution across the neuron population (i.e. the network state) evolves over time. If what we consider brain behavioral function depends on controlling its spatiotemporal patterns of neuron activity, we believe that this conceptual framework can be a useful starting point to analyze the properties of high-dimensional networks. The underlying representation obtained with this approach is the complete, high-dimensional vector field of the network, helping to preserve crucial information about its behavior that might otherwise be lost [19].

In order to get a more detailed understanding of the network state space and its dependence on connectivity and neuron activity, it is important to be able to visualize the vector field. However, as visualizing in high dimensions is in principle impossible, we introduced a planar subspace extraction tool for visualizing the combined action of a large number of neurons. This plane extraction tool can be used to illustrate any plane within the high dimensional network activity space. The extracted plane represents the interactions between two groups of neurons (one per axis), where a group could be a single neuron or, alternatively, any linear combination of neurons in the network (Fig 3). The choice of neuron groups is what defines the orientation of the extracted plane within the high-dimensional network activity space.

We showed that the presence of both excitatory and inhibitory neurons, with matching weights, makes it possible to displace the critical point from the origin (Fig 4) and to reshape the structure of the vector field around it. In this way, we can avoid obtaining networks which will inevitably be driven towards either zero neuron activity or towards saturation. A wider range of variations of the critical point location and its surrounding vector field could be obtained by skewing the synaptic weights (perhaps even including zero weights) (Fig 5). Another factor that can cause variations in the structure of the vector field is the neuronal activation function (Fig 7). Whereas the latter factors typically do not change to a great degree in the short term, any change in neuron activity will immediately impact the location of the critical point and the structure of the vector field that the rest of the neuron population ‘sees’ (‘perpendicular’ versus ‘non-perpendicular’ neurons in Fig 6). Since the impact of the vector field is to push the network towards new activity states, it should be possible to design such excitatory-inhibitory networks where the activity of the network constantly changes into new patterns, as in the brain *in vivo* [5,6,32,33]. As a consequence, the network, through the evolving vector field, would constantly be ‘chasing its own tail’, i.e. pushing itself towards a neutral point but never being able to reach it as the population activity would change as a consequence of the chasing process.

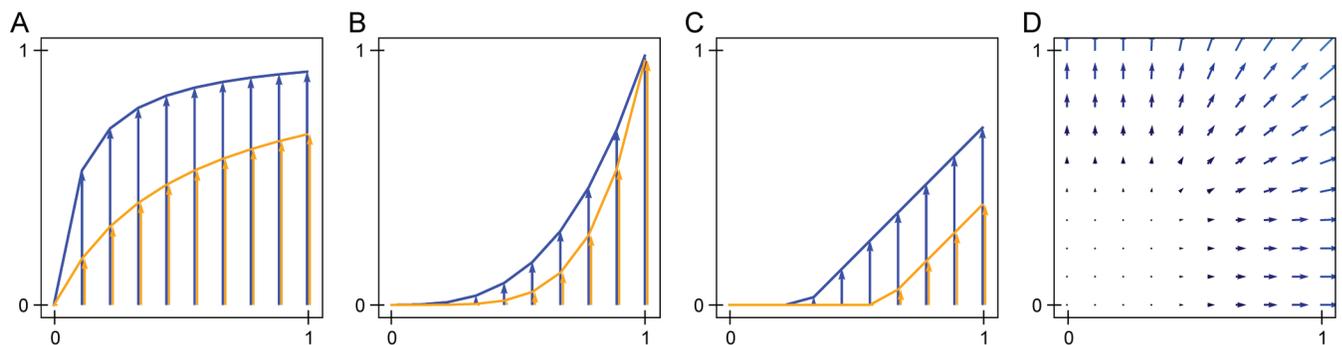


Fig 7. Impact of different neuronal input-output functions (f-I curves) on the vector fields. The impact of a neuron (i.e. the magnitude of its output vector component) in response to a given input is determined by the value of its f-I curve at that input level. (A) Output vectors for a decelerating activation function. The X axis represent the level of (summed synaptic) input that the neuron receives. The Y-axis represents the magnitude of the output vector, indicating the influence of the neuron on the activity of all downstream neurons receiving its synaptic input. This impact is subsequently scaled by the respective synaptic weights. The LSM used in the present study is an example of this type of activation function. The diagram illustrates the activation functions and the corresponding output vector components that result from two different levels of the static leak (blue: 0.1, orange: 0.5). (B) Same as in A but for an exponentially increasing, accelerating, activation function, here exemplified by the function $f(x) = x^\alpha$ with two settings of α (blue: 3, orange: 5). (C) Impact of two different thresholds on a neuron model with a simple linear activation function. An increase in the threshold results in that the vector components are shifted towards higher level of input activity but otherwise simply scale in proportion to the input. The threshold-linear response function $f(x) = \{0 \text{ if } x < \text{threshold}, (x - \text{threshold}) \text{ otherwise}\}$ was used with two threshold settings (blue: 0.3; orange: 0.6). (D) Vector field of a 2-excitatory neuron network (as in Fig 1B, 1C) with threshold-linear f-I curves with a threshold of 0.3. Vectors are scaled down for visibility.

<https://doi.org/10.1371/journal.pcsy.0000047.g007>

An implication of the vector field analogy is that, in the brain, the constantly evolving activity distribution across the neuron population will be a consequence of the synaptic connectivity landscape of the network. These 'trajectories' through the state space are the solutions provided by the network to given sensory inputs, at the cortical level [5,6] as well as in the spinal cord circuitry [8,13,14] and control how the activity distributions of the neuron population will evolve. The activity distributions, in turn, translate to, for example, the order and magnitude with which specific muscles are activated. Muscle activation patterns serve as the final expression of behavioral choices, making them a concrete example of how vector fields directly contribute to understanding the control of behavior.

According to the neuron model, the impact that a given individual presynaptic neuron can have on the vector field will be comparatively small as long as the activity of the other presynaptic neurons are non-zero and/or as long as the number of other presynaptic connections to the postsynaptic neuron is high (due to scaling of the static leak, see [Methods](#)). This implies that transient activations in single neurons cannot momentarily impact the population activity to a great extent. Since behavior ultimately consists in the spatiotemporal pattern of muscle activation, a behavioral change requires a coordinated change of the activity in larger populations of neurons to occur. However, provided that changes in the network population activity can be allowed to accumulate over some latency time, it is still likely that transient activity in individual neurons could form a seed to initiate a chain of activity changes in downstream neurons (see [34]). If this chain gradually moves the network in a new direction across the vector field because of that transient activation, for example during an ongoing movement, then it has contributed to the behavioral choice.

Notably there are many interesting implementations of modelled neural networks that operate as dynamical systems [35,36]. Our aim was not to design another variant of such a dynamical system but to introduce a conceptual framework that can be a useful starting point to analyze the population-level properties of high-dimensional networks, potentially bringing a better understanding of how brain networks may behave in this regard.

Limitations

In the examples provided throughout this paper, we illustrated vector fields at given levels of neuron output activities. Implicitly, we used the rate code of the neuronal output as a proxy for those activity levels. This is similar to the design of many studies that are interested in characterizing multi-neuronal interactions (Reviewed by [15,16], for specific examples also see [32,34,37]). It could be argued that neurons instead use discrete spikes in their communication. On the other hand, it has also been argued that the exact spike timing is not a deterministic process but a stochastic one, which is an argument for why a rate code approximation is a practical simplification [28,38–40]. Nevertheless, our vector field approach is also applicable for spiking neuronal networks, but then becomes a bit more involved to implement. Each spike would then need to be simulated to impact the output activity of the receiving neuron through the resulting synaptic potential, a potential that typically has a duration of many 10s of ms [41,42], although gradually declining. If the interspike intervals of the output neuron are shorter than the duration of the synaptic potential, then subsequent synaptic potentials will fuse and summate in a time-dependent manner, thus beginning to resemble the rate code we used here. A similar approach would be needed to emulate the impact of metabotropic synapses, where the membrane potential response to synaptic activation can last for more than 1000 ms [43]. To the extent that the metabotropic synapses regulate the gain of the ionotropic synapses [44], the vector field approach would then need to be

complemented with corresponding updates of the relevant vectors, i.e. the ionotropic synapses made on the neuron that receives the metabotropic input.

Short-term facilitation and depression of synaptic transmission at individual synapses was not included in this investigation, but could be modeled by scaling consecutive synaptic potentials, i.e. their template signals. This would be equivalent to dynamically adjusting synaptic weights, meaning that such short-term synaptic modifications could alternatively be captured by updating the vector field in a time-dependent manner. However, our vector field approach is intended to draw more attention to population-wide neuronal dynamics as an alternative to low-dimensional embeddings. When the population size is sufficiently large, highly specific details about individual neuronal and synaptic functions may have a tendency to average out when considering the overall neuron population behavior and dimensionality. Nevertheless, examining to what extent that applies remains a subject for future studies.

Future work

In neural recordings, the synaptic weights between the recorded neurons are typically unknown. Could estimations of the vector fields in the network still be extracted from recording data? Similar to the vector field applications for monitoring the spatial propagation of activity in the 3D tissue of the heart [9,10], multi-neuronal recordings, or even multi-location field potential recordings such as human EEG recordings, could be used to extract approximative vector fields in N-dimensional activity space. This approximation would be based on activity correlations between the neurons (or recording channels). Pairwise correlations between two neurons could be one step in this direction [45], but likely a better approach would be to use the changes in the activity distribution across the neurons for each consecutive time step to obtain an estimate of the vector at each location in the activity state space. This will be a subject for future work.

Conclusion

We used a simple neuron model to investigate how constraints on the neuron population level behavior would inevitably emerge from their synaptic interactions. We found that the distribution of synaptic weights greatly influences the structure of a vector field, while the activity distribution across the neuron population - i.e., the network's position within its activity state space — determines a vector which will influence how the network activity will change over time, thereby shaping how the activity distribution across the network evolves. The synaptic weight distributions are a primary factor driving the evolution of the activity distribution - which, in turn, is one of the most critical elements in defining brain function and behavior.

Supporting information

S1 File. Simulated neuronal network settings. The configuration settings of the neural network for each figure, including the locations of critical points where relevant.
(PY)

Author contributions

Conceptualization: Szilvia Szeier, Henrik Jörntell.

Formal analysis: Szilvia Szeier.

Investigation: Szilvia Szeier.

Methodology: Szilvia Szeier.

Project administration: Henrik Jörntell.

Software: Szilvia Szeier.

Supervision: Henrik Jörntell.

Validation: Szilvia Szeier.

Visualization: Szilvia Szeier.

Writing – original draft: Szilvia Szeier, Henrik Jörntell.

Writing – review & editing: Szilvia Szeier, Henrik Jörntell.

References

1. Douglas R, Koch C, Mahowald M, Martin K. The role of recurrent excitation in neocortical circuits. In: Models of cortical circuits. Springer; 1999. p. 251–82.
2. Binzegger T, Douglas RJ, Martin KAC. A quantitative map of the circuit of cat primary visual cortex. *J Neurosci*. 2004;24(39):8441–53. <https://doi.org/10.1523/JNEUROSCI.1400-04.2004> PMID: 15456817
3. Kalisman N, Silberberg G, Markram H. The neocortical microcircuit as a tabula rasa. *Proc Natl Acad Sci U S A*. 2005;102(3):880–5. <https://doi.org/10.1073/pnas.0407088102> PMID: 15630093
4. Enander JM, Spanne A, Mazzoni A, Bengtsson F, Oddo CM, Jörntell H. Ubiquitous neocortical decoding of tactile input patterns. *Front Cell Neurosci*. 2019;13:140.
5. Etemadi L, Enander JMD, Jörntell H. Remote cortical perturbation dynamically changes the network solutions to given tactile inputs in neocortical neurons. *iScience*. 2021;25(1):103557. <https://doi.org/10.1016/j.isci.2021.103557> PMID: 34977509
6. Norrliid J, Enander JMD, Mogensen H, Jörntell H. Multi-structure cortical states deduced from intracellular representations of fixed tactile input patterns. *Front Cell Neurosci*. 2021;15:677568. <https://doi.org/10.3389/fncel.2021.677568> PMID: 34194301
7. Wahlbom A, Enander JMD, Bengtsson F, Jörntell H. Focal neocortical lesions impair distant neuronal information processing. *J Physiol*. 2019;597(16):4357–71. <https://doi.org/10.1113/JP277717> PMID: 31342538
8. Kohler M, Bengtsson F, Stratmann P, Röhrbein F, Knoll A, Albu-Schäffer A, et al. Diversified physiological sensory input connectivity questions the existence of distinct classes of spinal interneurons. *iScience*. 2022;25(4):104083. <https://doi.org/10.1016/j.isci.2022.104083> PMID: 35372805
9. Dallet C, Roney C, Martin R, Kitamura T, Puyo S, Duchateau J, et al. Cardiac propagation pattern mapping with vector field for helping tachyarrhythmias diagnosis with clinical tridimensional electro-anatomical mapping tools. *IEEE Trans Biomed Eng*. 2019;66(2):373–82. <https://doi.org/10.1109/TBME.2018.2841340> PMID: 29993411
10. Pancorbo L, Ruipérez-Campillo S, Tormos Á, Guill A, Cervigón R, Alberola A. Vector field heterogeneity for the assessment of locally disorganised cardiac electrical propagation wavefronts from high-density multielectrodes. *IEEE Open J Eng Med Biol*. 2023.
11. Townsend RG, Gong P. Detection and analysis of spatiotemporal patterns in brain activity. *PLoS Comput Biol*. 2018;14(12):e1006643. <https://doi.org/10.1371/journal.pcbi.1006643> PMID: 30507937
12. Yu BM, Kemere C, Santhanam G, Afshar A, Ryu SI, Meng TH, et al. Mixture of trajectory models for neural decoding of goal-directed movements. *J Neurophysiol*. 2007;97(5):3763–80. <https://doi.org/10.1152/jn.00482.2006> PMID: 17329627
13. Kohler M, Röhrbein F, Knoll A, Albu-Schäffer A, Jörntell H. The Bcm rule allows a spinal cord model to learn rhythmic movements. *Biol Cybern*. 2023;117(4–5):275–84. <https://doi.org/10.1007/s00422-023-00970-z> PMID: 37594531
14. Santello M, Baud-Bovy G, Jörntell H. Neural bases of hand synergies. *Front Comput Neurosci*. 2013;7:23. <https://doi.org/10.3389/fncom.2013.00023> PMID: 23579545
15. Gallego JA, Perich MG, Miller LE, Solla SA. Neural manifolds for the control of movement. *Neuron*. 2017;94(5):978–84. <https://doi.org/10.1016/j.neuron.2017.05.025> PMID: 28595054

16. Khona M, Fiete IR. Attractor and integrator networks in the brain. *Nat Rev Neurosci*. 2022;23(12):744–66. <https://doi.org/10.1038/s41583-022-00642-0> PMID: 36329249
17. Luczak A, Barthó P, Harris KD. Spontaneous events outline the realm of possible sensory responses in neocortical populations. *Neuron*. 2009;62(3):413–25. <https://doi.org/10.1016/j.neuron.2009.03.014> PMID: 19447096
18. Jazayeri M, Ostojic S. Interpreting neural computations by examining intrinsic and embedding dimensionality of neural activity. *Curr Opin Neurobiol*. 2021;70:113–20. <https://doi.org/10.1016/j.conb.2021.08.002> PMID: 34537579
19. Kristensen SS, Kesgin K, Jörntell H. High-dimensional cortical signals reveal rich bimodal and working memory-like representations among S1 neuron populations. *Commun Biol*. 2024;7(1):1043. <https://doi.org/10.1038/s42003-024-06743-z> PMID: 39179675
20. Pellegrino A, Stein H, Cayco-Gajic NA. Dimensionality reduction beyond neural subspaces with slice tensor component analysis. *Nat Neurosci*. 2024;27(6):1199–210. <https://doi.org/10.1038/s41593-024-01626-2> PMID: 38710876
21. Musall S, Kaufman MT, Juavinett AL, Gluf S, Churchland AK. Single-trial neural dynamics are dominated by richly varied movements. *Nat Neurosci*. 2019;22(10):1677–86. <https://doi.org/10.1038/s41593-019-0502-4> PMID: 31551604
22. Rongala UB, Enander JMD, Kohler M, Loeb GE, Jörntell H. A non-spiking neuron model with dynamic leak to avoid instability in recurrent networks. *Front Comput Neurosci*. 2021;15:656401. <https://doi.org/10.3389/fncom.2021.656401> PMID: 34093156
23. Hodgkin AL, Huxley AF. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol*. 1952;117(4):500–44. <https://doi.org/10.1113/jphysiol.1952.sp004764> PMID: 12991237
24. Spruston N, Johnston D. Perforated patch-clamp analysis of the passive membrane properties of three classes of hippocampal neurons. *J Neurophysiol*. 1992;67(3):508–29. <https://doi.org/10.1152/jn.1992.67.3.508> PMID: 1578242
25. Thurbon D, Lüscher HR, Hofstetter T, Redman SJ. Passive electrical properties of ventral horn neurons in rat spinal cord slices. *J Neurophysiol*. 1998;79(5):2485–502. <https://doi.org/10.1152/jn.1998.79.5.2485> PMID: 9722433
26. Philippens IH. Marmosets in neurologic disease research: Parkinson's disease. The common marmoset in captivity and biomedical research. Amsterdam: Elsevier. 2019. p. 415–35.
27. Pavlov I, Savtchenko LP, Kullmann DM, Semyanov A, Walker MC. Outwardly rectifying tonically active GABAA receptors in pyramidal cells modulate neuronal offset, not gain. *J Neurosci*. 2009;29(48):15341–50. <https://doi.org/10.1523/JNEUROSCI.2747-09.2009> PMID: 19955387
28. Spanne A, Geborek P, Bengtsson F, Jörntell H. Spike generation estimated from stationary spike trains in a variety of neurons in vivo. *Front Cell Neurosci*. 2014;8:199. <https://doi.org/10.3389/fncel.2014.00199> PMID: 25120429
29. Goldberg EM, Clark BD, Zagha E, Nahmani M, Erisir A, Rudy B. K⁺ channels at the axon initial segment dampen near-threshold excitability of neocortical fast-spiking GABAergic interneurons. *Neuron*. 2008;58(3):387–400. <https://doi.org/10.1016/j.neuron.2008.03.003> PMID: 18466749
30. Azouz R, Gray CM. Dynamic spike threshold reveals a mechanism for synaptic coincidence detection in cortical neurons in vivo. *Proc Natl Acad Sci U S A*. 2000;97(14):8110–5. <https://doi.org/10.1073/pnas.130200797> PMID: 10859358
31. de Polavieja GG, Harsch A, Kleppe I, Robinson HPC, Juusola M. Stimulus history reliably shapes action potential waveforms of cortical neurons. *J Neurosci*. 2005;25(23):5657–65. <https://doi.org/10.1523/JNEUROSCI.0242-05.2005> PMID: 15944394
32. Stringer C, Pachitariu M, Steinmetz N, Reddy CB, Carandini M, Harris KD. Spontaneous behaviors drive multidimensional, brainwide activity. *Science*. 2019;364(6437):255. <https://doi.org/10.1126/science.aav7893> PMID: 31000656
33. Nguyen ND, Lutas A, Amsalem O, Fernando J, Ahn AY-E, Hakim R, et al. Cortical reactivations predict future sensory responses. *Nature*. 2024;625(7993):110–8. <https://doi.org/10.1038/s41586-023-06810-1> PMID: 38093002
34. Allen WE, Chen MZ, Pichamoorthy N, Tien RH, Pachitariu M, Luo L, et al. Thirst regulates motivated behavior through modulation of brainwide neural population dynamics. *Science*. 2019;364(6437):253. <https://doi.org/10.1126/science.aav3932> PMID: 30948440
35. Boerlin M, Machens CK, Denève S. Predictive coding of dynamical variables in balanced spiking networks. *PLoS Comput Biol*. 2013;9(11):e1003258. <https://doi.org/10.1371/journal.pcbi.1003258> PMID: 24244113
36. Podlaski WF, Machens CK. Approximating nonlinear functions with latent boundaries in low-rank excitatory-inhibitory spiking networks. *Neural Comput*. 2024;36(5):803–57. https://doi.org/10.1162/neco_a_01658 PMID: 38658028

37. Bimbard C, Sit TPH, Lebedeva A, Reddy CB, Harris KD, Carandini M. Behavioral origin of sound-evoked activity in mouse visual cortex. *Nat Neurosci*. 2023;26(2):251–8. <https://doi.org/10.1038/s41593-022-01227-x> PMID: 36624279
38. Naundorf B, Wolf F, Volgushev M. Unique features of action potential initiation in cortical neurons. *Nature*. 2006;440(7087):1060–3. <https://doi.org/10.1038/nature04610> PMID: 16625198
39. Saarinen A, Linne M-L, Yli-Harja O. Stochastic differential equation model for cerebellar granule cell excitability. *PLoS Comput Biol*. 2008;4(2):e1000004. <https://doi.org/10.1371/journal.pcbi.1000004> PMID: 18463700
40. Nilsson MNP, Jörntell H. Channel current fluctuations conclusively explain neuronal encoding of internal potential into spike trains. *Phys Rev E*. 2021;103(2–1):022407. <https://doi.org/10.1103/PhysRevE.103.022407> PMID: 33736029
41. Reyes A, Sakmann B. Developmental switch in the short-term modification of unitary EPSPs evoked in layer 2/3 and layer 5 pyramidal neurons of rat neocortex. *J Neurosci*. 1999;19(10):3827–35. <https://doi.org/10.1523/JNEUROSCI.19-10-03827.1999> PMID: 10234015
42. Lefort S, Tomm C, Floyd Sarria J-C, Petersen CCH. The excitatory neuronal network of the C2 barrel column in mouse primary somatosensory cortex. *Neuron*. 2009;61(2):301–16. <https://doi.org/10.1016/j.neuron.2008.12.020> PMID: 19186171
43. Batchelor AM, Garthwaite J. Frequency detection and temporally dispersed synaptic signal association through a metabotropic glutamate receptor pathway. *Nature*. 1997;385(6611):74–7. <https://doi.org/10.1038/385074a0> PMID: 8985249
44. Stratmann P, Albu-Schäffer A, Jörntell H. Scaling our world view: how monoamines can put context into brain circuitry. *Front Cell Neurosci*. 2018;12:506. <https://doi.org/10.3389/fncel.2018.00506> PMID: 30618646
45. Cohen MR, Kohn A. Measuring and interpreting neuronal correlations. *Nat Neurosci*. 2011;14(7):811–9. <https://doi.org/10.1038/nn.2842> PMID: 21709677