

RESEARCH ARTICLE

Regimes and mechanisms of transient amplification in abstract and biological neural networks

Georgia Christodoulou¹, Tim P. Vogels^{1,2}, Everton J. Agnes^{1,3}***1** Centre for Neural Circuits and Behaviour, University of Oxford, Oxford, United Kingdom, **2** Institute of Science and Technology Austria, Klosterneuburg, Austria, **3** Biozentrum, University of Basel, Basel, Switzerland

* These authors contributed equally to this work.

* everton.agnes@gmail.com

Abstract

Neuronal networks encode information through patterns of activity that define the networks' function. The neurons' activity relies on specific connectivity structures, yet the link between structure and function is not fully understood. Here, we tackle this structure-function problem with a new conceptual approach. Instead of manipulating the connectivity directly, we focus on upper triangular matrices, which represent the network dynamics in a given orthonormal basis obtained by the Schur decomposition. This abstraction allows us to independently manipulate the eigenspectrum and feedforward structures of a connectivity matrix. Using this method, we describe a diverse repertoire of non-normal transient amplification, and to complement the analysis of the dynamical regimes, we quantify the geometry of output trajectories through the effective rank of both the eigenvector and the dynamics matrices. Counter-intuitively, we find that shrinking the eigenspectrum's imaginary distribution leads to highly amplifying regimes in linear and long-lasting dynamics in nonlinear networks. We also find a trade-off between amplification and dimensionality of neuronal dynamics, i.e., trajectories in neuronal state-space. Networks that can amplify a large number of orthogonal initial conditions produce neuronal trajectories that lie in the same subspace of the neuronal state-space. Finally, we examine networks of excitatory and inhibitory neurons. We find that the strength of global inhibition is directly linked with the amplitude of amplification, such that weakening inhibitory weights also decreases amplification, and that the eigenspectrum's imaginary distribution grows with an increase in the ratio between excitatory-to-inhibitory and excitatory-to-excitatory connectivity strengths. Consequently, the strength of global inhibition reveals itself as a strong signature for amplification and a potential control mechanism to switch dynamical regimes. Our results shed a light on how biological networks, i.e., networks constrained by Dale's law, may be optimised for specific dynamical regimes.

OPEN ACCESS

Citation: Christodoulou G, Vogels TP, Agnes EJ (2022) Regimes and mechanisms of transient amplification in abstract and biological neural networks. *PLoS Comput Biol* 18(8): e1010365. <https://doi.org/10.1371/journal.pcbi.1010365>

Editor: Yashar Ahmadian, University of Oregon, UNITED STATES

Received: June 17, 2021

Accepted: July 6, 2022

Published: August 15, 2022

Copyright: © 2022 Christodoulou et al. 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: Relevant computer code is available at the GitHub repository https://github.com/ejagnes/transient_amplification_abst_bio/.

Funding: This research was funded by a Wellcome Trust and Royal Society Henry Dale Research Fellowship (WT100000; TPV), a Wellcome Senior Research Fellowship (214316/Z/18/Z; GC, EJA, and TPV), and a Research Project Grant by the Leverhulme Trust (RPG-2016-446; EJA and TPV). 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.

Author summary

The architecture of neuronal networks lies at the heart of its dynamic behaviour, or in other words, the function of the system. However, the relationship between changes in the architecture and their effect on the dynamics, a structure-function problem, is still poorly understood. Here, we approach this problem by studying a rotated connectivity matrix that is easier to manipulate and interpret. We focus our analysis on a dynamical regime that arises from the biological property that neurons are usually not connected symmetrically, which may result in a non-normal connectivity matrix. Our techniques unveil distinct expressions of the dynamical regime of non-normal amplification. Moreover, we devise a way to analyse the geometry of the dynamics: we assign a single number to a network that quantifies how dissimilar its repertoire of behaviours can be. Finally, using our approach, we can close the loop back to the original neuronal architecture and find that biologically plausible networks use the strength of inhibition and excitatory-to-inhibitory connectivity strength to navigate the different dynamical regimes of non-normal amplification.

Introduction

Recurrent network models are known to produce different types of dynamics, ranging from regular to irregular, and from transient to persistent activity [1–6]. Moulding network dynamics to resemble experimental observations usually involves changes in the network architecture, i.e., the existence of synapses and their efficacies [7–9]. With this approach, the eigenspectrum and the non-normality of the connectivity matrix are indirectly affected, and the relationship between changes in those qualities of the weight matrix and the network dynamics remain unclear. This is challenging because both the eigenspectrum and non-normality carry important information about the dynamics. The eigenspectrum, i.e., the distribution of eigenvalues in the complex plane, carries information about the stability of the network (e.g., asymptotic behaviour) [6, 8, 10], and timescale of the dynamics [4, 11]. However, the eigenspectrum alone is not sufficient to describe the transient dynamical behaviour of a network [10, 12]. The transient dynamics, and more specifically, the phenomenon of transient amplification, depends on the alignment between the eigenvectors of the connectivity matrix [10]. Importantly, the more aligned the eigenvectors are, the more *non-normal* a matrix is. The non-normality of a matrix can be assessed through the Schur decomposition, an orthogonal similarity transformation that results in an upper triangular matrix on which the eigenvalues appear along the diagonal. The Schur decomposition of a matrix \mathbf{W} can be formally written as $\mathbf{W} = \mathbf{U}(\mathbf{\Lambda} + \mathbf{T})\mathbf{U}^\dagger$, where \mathbf{U} is a unitary matrix (its columns are the orthogonal Schur modes), $\mathbf{\Lambda}$ is a diagonal matrix containing the eigenvalues of \mathbf{W} (complex eigenvalues are represented as 2-by-2 block diagonals with the real and imaginary parts), and \mathbf{T} is a strictly upper triangular matrix. The strictly upper triangular part contains information related to the interactions between the corresponding Schur modes [12]—it's useful to note that a normal matrix has this strictly upper triangular part equal to zero because all eigenvectors are orthogonal. As such, the strength (measured by the norm [13]) of the strictly triangular part of the Schur decomposition plays an important role for the dynamics.

Therefore here, we consider upper triangular matrices and manipulate their spectrum and non-normality, such that their characteristics can be directly translated into dynamical properties (Fig 1A). These matrices no longer represent the neuronal connectivity, but modes of activation that are arranged in a feedforward manner [10, 12, 14] (Fig 1B). We are particularly

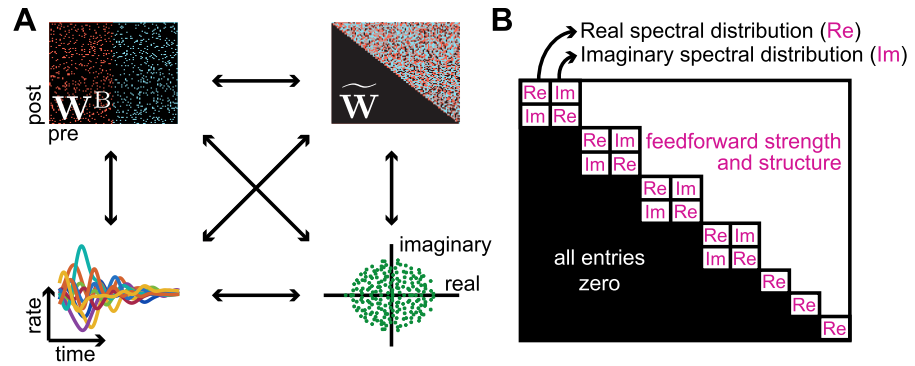


Fig 1. Elements explored in the manuscript. A, Schematic of the elements explored in this manuscript. Top left and clockwise: The connectivity matrix W^B ; its corresponding Schur upper triangular decomposition \tilde{W} ; the eigenspectrum; and the induced dynamics. In W^B and \tilde{W} , red dots indicate positive (excitatory) connections while blue dots indicate negative (inhibitory) connections. B, The upper triangular matrix \tilde{W} with the quantities that we alter in this manuscript in pink.

<https://doi.org/10.1371/journal.pcbi.1010365.g001>

interested in the different forms of transient amplification, a phenomenon that can resemble motor cortex activity during reaching [15–17] and also emulate long-lasting working memory dynamics [18–20]. For a network to be able to transiently amplify an input signal (i.e., initial condition), not only the eigenvectors need not be orthogonal (allowing for non-normal amplification), but also eigenvalues cannot have any real part larger than unity (allowing for the transient amplification to be followed by a decay back to baseline) [6, 10, 12]. This constrains the possible structure of the eigenspectrum. For example, when the connectivity matrix is random, the bulk of the eigenvalues is uniformly distributed on a disc centred at zero, which together with the stability constraint imposes that the radius can't be larger than 1 [2, 6, 12]. Taking into account this effect, random weight matrices are not particularly flexible for generating distinct eigenspectrum distributions, not allowing for a rich plethora of network dynamics. Specific eigenspectrum distributions (given that the matrix is non-normal), can indeed elicit richer dynamics through either optimisation of weights taking into consideration their effect on the eigenspectrum [8] or the combination of matrices with different statistics [6, 21, 22]. However, it is still unclear how modifications in the distribution of only real or imaginary parts of the eigenspectrum change the dynamical regimes of a non-normal network which is able to transiently amplify inputs. Moreover, it's not entirely known how these changes can be imposed to biologically realistic networks in which neurons are either excitatory or inhibitory, and neurons might not connect to themselves, i.e., without self connections.

To better understand how the manipulations in eigenspectrum and non-normality of a connectivity matrix are translated to biologically plausible networks, in this work we manipulate them directly in upper triangular matrices and then translate our main findings to biologically plausible networks. We start by defining the relationship between the network's activity and the eigenmodes given by the Schur decomposition, showing how simple manipulations of the eigenspectrum's imaginary diameter and real distribution can drastically affect the neuronal dynamics. We then systematically explore how the eigenspectrum's imaginary diameter and the feedforward norm control the different regimes of transient amplification, showing that either shrinking the eigenspectrum's imaginary diameter or increasing the feedforward norm increase amplification levels. Analysing the neuronal dynamics of these networks via the effective rank of the eigenvectors of the connectivity matrix, we find a trade-off between amplification and the underlying dimensionality of the possible dynamics of the network: networks

with high (respectively low) levels of amplification produce dynamics in low (respectively high) dimensional subspaces. After a dissection of the underlying mechanisms of transient amplification using general upper triangular matrices, we consider biological constraints on the spectral distributions, and consequently, on the dynamics. Finally, we show how we can implement our findings in a biological plausible connectivity matrix with excitatory and inhibitory neurons, i.e., a matrix satisfying Dale’s law. We explore three different manipulations: lack of autapses, global inhibition, and ratio between excitatory-to-inhibitory and excitatory-to-excitatory connections. We show that increased global inhibition can lead to more amplifying dynamics due to its connections with the eigenspectrum, and that changing excitatory-to-inhibitory connections affect the eigenspectrum’s imaginary diameter.

Results

Throughout the paper we use the following notation for the connectivity matrix: \mathbf{W} for a generic connectivity matrix, $\widetilde{\mathbf{W}}$ for a matrix given in upper triangular form, and \mathbf{W}^B for a matrix following biological constraints (Fig 1A). The dynamics of the recurrent network are defined by

$$\tau \frac{d\mathbf{x}(t)}{dt} = -\mathbf{x}(t) + \mathbf{W}f(\mathbf{x}(t)), \tag{1}$$

where $\mathbf{x}(t)$ is the internal state of the network at time t , and $x_i(t)$ can be understood as the membrane potential of the i^{th} neuron of the network (with $i = 1, \dots, N$; N is the number of neurons in the network). This internal state of the neurons evolves with a characteristic time constant τ (fixed at $\tau = 200$ ms throughout the paper) and is affected by the activity of other neurons of the network through the recurrent connections determined by \mathbf{W} . Finally, the activation function, $f(\mathbf{x}(t)) = \mathbf{r}(t)$, represents the input-output relation between the internal state, $\mathbf{x}(t)$, and the firing rate deviation, $\mathbf{r}(t)$, from the baseline activity, \mathbf{r}_0 . We assume linear dynamics, $\mathbf{r} = f(\mathbf{x}) = \mathbf{x}$, for the mathematical analysis and compare to networks with richer dynamics using a known non-linear function given by

$$f(x) = \begin{cases} r_{\min} \tanh\left(\frac{x}{r_{\min}}\right) & \text{for } x < 0 \\ r_{\max} \tanh\left(\frac{x}{r_{\max}}\right) & \text{otherwise,} \end{cases} \tag{2}$$

where $r_{\min} = 1$ Hz and $r_{\max} = 4$ Hz are the bounds of the sigmoid function $f(x)$. Note that in this case, and generally used in previous works [4, 8], negative values of $r_i(t)$ means less than baseline activity (see S1 Text for a non-negative version). In the linear case, the network dynamics can be described using the eigenvalues, λ_i , and eigenvectors, \mathbf{v}_i , of the weight matrix \mathbf{W} . To quantify whether and by how much the network can amplify specific inputs, we calculate the norm of the rate vector, $\|\mathbf{r}(t)\|$, by decomposing it along the directions of the eigenvectors of \mathbf{W} ,

$$\|\mathbf{r}(t)\| = \sqrt{\sum_k |r_k(t)|^2 + \sum_{k \neq j} \overline{r_k(t)} r_j(t) \langle \mathbf{v}_k, \mathbf{v}_j \rangle}, \tag{3}$$

where $\overline{r_k(t)}$ is the complex conjugate of $r_k(t)$ and $\langle \mathbf{v}_k, \mathbf{v}_j \rangle$ is the inner product of the complex vectors \mathbf{v}_k and \mathbf{v}_j (see Methods). Here, $r_k(t) = \hat{r}_k \exp\left(\frac{t}{\tau}(\lambda_k - 1)\right)$ is the solution of the system along the direction of the eigenvector \mathbf{v}_k , which is associated with the eigenvalue λ_k (\hat{r}_k is a constant, uniquely determined by the initial condition). In a stable regime, $\text{Re}(\lambda_k) < 1, \forall k$, the

system exhibits a single fixed point that represents the baseline activity. An increase of the response norm, $\|\mathbf{r}(t)\|$, with respect to the norm of the initial condition, $\|\mathbf{r}(t_0)\|$ (here always normalised to 1), defines the phenomenon of transient amplification. A necessary condition for this to happen is the non-normality of \mathbf{W} , i.e., the eigenvectors do not form an orthogonal basis [10]: $\langle \mathbf{v}_k, \mathbf{v}_j \rangle \neq 0$, for some j, k .

To explore regimes of transient amplification, we thus focus on matrices of the form $\widetilde{\mathbf{W}} = \mathbf{\Lambda} + \mathbf{T}$ (Fig 1B), with the diagonal, $\mathbf{\Lambda}$, containing the eigenvalues [10, 12, 14], and the strictly upper triangular part, \mathbf{T} , representing the feedforward structure between patterns of activation (see S1 Text). Note that $\mathbf{\Lambda}$ contains 2×2 -blocks around the diagonal to accommodate for complex eigenvalues in real-valued matrices. The real parts of the eigenvalues are on the diagonal and the imaginary parts lie on the off-diagonal entries of the 2×2 blocks (Fig 1B; see Methods). We create $\mathbf{\Lambda}$ by sampling the real and imaginary parts of the eigenvalues from different distributions, but keeping the number of complex versus real eigenvalues constant (see Methods). The imaginary distribution needs to be symmetric with respect to zero (a condition imposed by the conjugacy of the complex eigenvalues), while the real distribution must be below 1 (and is here always set to have 0.5 as a supremum: $\max_k \text{Re}(\lambda) < 0.5$) to ensure stability. For our analysis, we define the *spectrum norm* as the Frobenius norm of the matrix $\mathbf{\Lambda}$,

$\|\mathbf{\Lambda}\| = \sqrt{\sum_k |\lambda_k|^2}$ (note that the spectrum norm defined here is different than the commonly used *spectral norm* [23]). We create \mathbf{T} in two different ways: from the Schur decomposition of a Stability-Optimised Circuit (SOC) [8] or sampled from a uniform distribution. A SOC is created from an initially unstable network (i.e., the initial weight matrix has some of the eigenvalues with real part greater than unity, $\text{Re}(\lambda_k) > 1$ for some k), whose inhibitory connections are modified (optimised) so that no eigenvalue has real part greater than unity (see Methods). After optimisation, a SOC produces strong non-normal transient amplification [8]. We thus use the SOC's corresponding feedforward structure because it is tuned to create transient amplification. We linearly scale all the elements of \mathbf{T} after its structure has been fixed to vary its norm.

The subsequent sections are organised as follows. First, we explore the effects of changing the distribution of the eigenspectrum's imaginary and real parts. We then examine the influence of the spectrum and feedforward norms, $\|\mathbf{\Lambda}\|$ and $\|\mathbf{T}\|$, respectively, for different regimes of transient amplification, analysing the dimensionality of the dynamics in such regimes. Finally, we link some of the findings from abstract to biological networks through manipulations in a biologically realistic weight matrix that satisfies Dale's law.

Shrinking the imaginary distribution increases amplification

We start our investigation of how the eigenspectrum affects the dynamics by drawing both real and imaginary parts from uniform distributions with diameters d_{im} and d_{re} , respectively (Fig 2A–2D, top left). To quantify the dynamical response of the network, we find an orthogonal basis of initial conditions of the linear network that elicit maximum amplification, $\mathcal{I}_B = \{\mathbf{a}_1, \dots, \mathbf{a}_N\}$, ordered according to their *evoked energy* [8], $\mathcal{E}(\mathbf{a})$: if $\mathbf{x}(t=0) = \mathbf{a}$, $\|\mathbf{a}\| = 1$, then $\mathcal{E}(\mathbf{a}) = \frac{2}{\tau} \int_0^\infty \|\mathbf{x}(t)\|^2 dt$. We first calculate the initial condition \mathbf{a}_1 that maximises $\mathcal{E}(\mathbf{a})$, following an iterative calculation of the subsequent orthogonal initial condition, \mathbf{a}_i , that maximise $\mathcal{E}(\mathbf{a})$ in the subspace orthogonal to the previously calculated initial conditions, $\{\mathbf{a}_1, \dots, \mathbf{a}_{i-1}\}$. To make sure that the evoked energy is due to an amplified response rather than merely a slower exponential decay, we compute the maximum value of the norm of the firing rate vector, for all vectors in \mathcal{I}_B (Fig 2A–2D, bottom left).

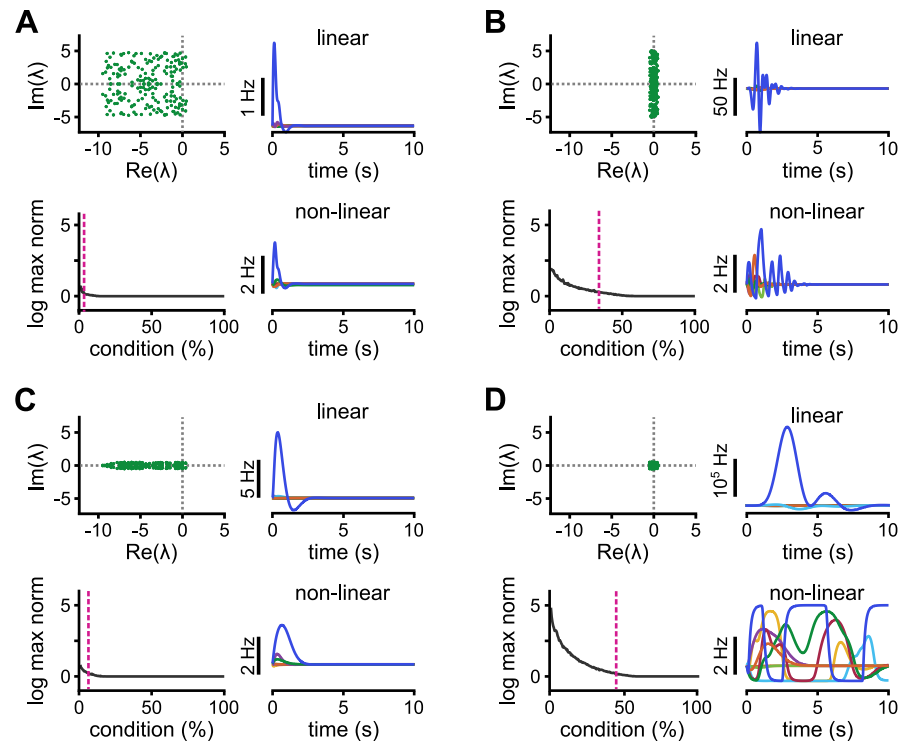


Fig 2. Eigenspectra and dynamics of the corresponding networks. A–D, Four cases of eigenspectra and dynamics of the corresponding network of size $N = 200$. In each panel, clockwise: The spectrum; linear dynamics; non-linear dynamics; the logarithm of the maximum norm of the firing rate per initial condition. The same initial condition that elicits the maximum norm is used for both linear and non-linear dynamics. Pink dotted line indicates the percentage of conditions whose norm is amplified by at least 50%. The feedforward structure is taken from a stability-optimised circuit [8] and its Frobenius norm is fixed to 75. Real and imaginary parts follow an uniform distribution with diameters d_{im} and d_{re} , respectively. **A**, When $d_{im} = d_{re} = 10$, only 1% (2 out of 200) of the conditions are slightly amplified. **B**, When $d_{im} = 10$ and $d_{re} = 1$, the system is capable of more amplification. **C**, Here, $d_{im} = 1$ and $d_{re} = 10$, surprisingly creating more amplification compared to the case shown in panel A. **D**, When $d_{im} = d_{re} = 1$, the system amplifies almost half of the initial conditions. The dynamics, given an initial condition of norm 1, reach the value of $\sim 10^5$ Hz in the linear case and consequently long-lasting dynamics in the non-linear case.

<https://doi.org/10.1371/journal.pcbi.1010365.g002>

With broad distributions, the system can slightly amplify a few conditions (Fig 2A). When the range of the real-part distribution is decreased and the real parts are pushed towards 0.5, the resulting network produces stronger amplification (Fig 2B). This can mainly be attributed to the fact that the eigenvalues have now larger real parts and hence longer decay envelopes. A longer decay envelope allows for the network dynamics to undergo a larger proportion of a full period of the underlying oscillation without damping of the maximum amplitude, revealing the hidden feedforward structure, and thus increasing the maximum norm. Indeed, clustering away from 0.5 leads to less amplification (explored in the next section).

More surprisingly, shrinking, instead, the imaginary distribution also leads to more amplification (Fig 2C), and shrinking both distributions produces very large amplification that in the non-linear case lasts for a long time (longer than 10 seconds), approximating timescales of working memory dynamics (Fig 2D), previously known to arise through spectral abscissas near the stability line. Additionally, the percentage of conditions that are amplified is considerably increased, i.e., the ability of such a network to amplify orthogonal initial conditions is enhanced. Note that splitting and clustering the (positive and negative) imaginary parts away from zero gives rise to slightly different amplification regimes that also depend on the linearity of the system (S1 Fig). From these examples (Fig 2 and S1 Fig) it is clear that varying the

diameter and position of the centre of the distribution of both imaginary and real parts of the eigenspectrum play distinct roles in the levels of amplification of a network undergoing transient amplification.

Manipulating the spectrum and the feedforward norms

When we study the effects of the imaginary and real distributions more systematically, we find that the shape of the real distributions affects the levels of amplification, but has minimum effect on how amplification changes with the imaginary distribution (Fig 3Ai and 3Aii and S2(A) Fig). Amplification emerges from the non-normality of W , which can be partly

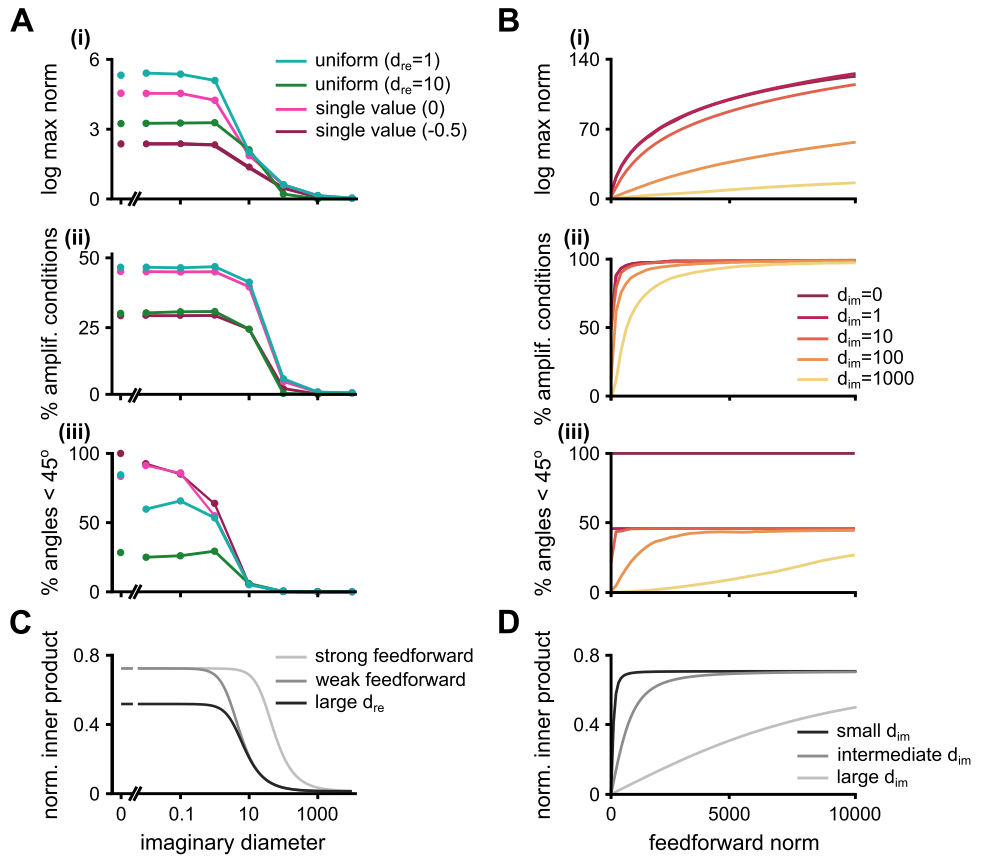


Fig 3. Effects of manipulating the spectrum and the feedforward norms. **A**, Maximum response norm for the preferred initial condition (i), percentage of directions whose norm is amplified more than 50% (ii), and the percentage of angles (between pairs of eigenvectors) that are less than 45° (iii). Every line is a function of the imaginary diameter. We plot four real distributions. Light green: a uniform distribution in which all real parts are distributed uniformly in the interval $(-0.5, 0.5)$. Dark green: a uniform distribution in which all real parts are distributed uniformly in the interval $(-9.5, 0.5)$. Light pink: a single valued real distribution in which all real parts are equal to zero. Dark pink: a single valued real distribution in which all real parts are equal to -0.5 . In all cases the network size is $N = 200$ and the feedforward Frobenius norm is fixed at 75. d_{re} indicates the diameter of the uniform distribution of the eigenspectrum's real part. **B**, Same as panel A, but plotted as a function of the feedforward Frobenius norm. Different colours correspond to 5 different spectra; all spectra have fixed single-valued real distributions (equal to zero) and different imaginary diameters. d_{im} indicates the diameter of the uniform distribution of the eigenspectrum's imaginary part. **C**, Normalised inner product between vectors from a simplified 3-by-3 upper triangular matrix (Eq 4) as a function of the imaginary diameter (β in Eq 4) for three conditions: "strong feedforward" ($\phi_{norm} = 30$ and $\alpha - \gamma = -0.3$); "weak feedforward" ($\phi_{norm} = 3$ and $\alpha - \gamma = -0.3$); and "large d_{re} " ($\phi_{norm} = 3$ and $\alpha - \gamma = -3$). **D**, Same as panel C, but plotted as a function of the feedforward norm for three different conditions: "small d_{im} " ($\beta = 100$ and $\alpha - \gamma = 0$); "intermediate d_{im} " ($\beta = 1000$ and $\alpha - \gamma = 0$); and "large d_{im} " ($\beta = 10000$ and $\alpha - \gamma = 0$).

<https://doi.org/10.1371/journal.pcbi.1010365.g003>

quantified by the angles between the eigenvectors (Eq 3); if more pairs have overlaps, the matrix will be more non-normal. The imaginary distribution changes the geometry of the eigenvectors (Fig 3Aiii), providing a mechanism for its drastic effect on the amplification in these networks (Fig 3Ai and 3Aii). This is a surprising effect given that we do not alter the feedforward norm, $\|\mathbf{T}\|$, i.e., the Frobenius norm of the strictly upper triangular part of $\widetilde{\mathbf{W}}$, nor the decay envelopes at all (see S2(B) Fig and S1 Text for variations of feedforward structures).

The feedforward norm is more directly linked to the non-normality [10], and as expected, it increases both the norm of the maximum response (Fig 3Bi), and the percentage of amplified conditions (Fig 3Bii), for larger values. The percentage of eigenvector pairs with small angles also grows with increasing feedforward norm strength (Fig 3Biii). Interestingly, there is a saturating point of eigenvector pairs aligned with angles smaller than 45° that depends on the imaginary distribution. Once the number of pairs saturates, increased amplification may be associated with an increase in the matrix norm, $\|\mathbf{W}\|$, or further alignment of these eigenvector pairs (even smaller angles).

To get an intuition of the mechanisms behind the changes in amplification levels we analysed a 3-by-3 upper triangular matrix with one purely real eigenvalue, $\lambda_1 = \gamma$, and two complex eigenvalues, $\lambda_2 = \overline{\lambda_3} = \alpha - i\beta$. Two out of the three eigenvectors of this matrix are orthogonal, $\langle \mathbf{v}_2, \mathbf{v}_3 \rangle = 0$. The inner product of the non-orthogonal eigenvectors is given by (see Methods for details)

$$\frac{|\langle \mathbf{v}_1, \mathbf{v}_2 \rangle|}{\|\mathbf{v}_1\| \|\mathbf{v}_2\|} = \frac{|\langle \mathbf{v}_1, \mathbf{v}_3 \rangle|}{\|\mathbf{v}_1\| \|\mathbf{v}_3\|} = \frac{\phi_{\text{norm}}}{\sqrt{2[\phi_{\text{norm}}^2 + (\alpha - \gamma)^2 + \beta^2]}}, \quad (4)$$

where ϕ_{norm} is the feedforward norm. The result from the simplified 3-by-3 matrix follows the same trends as the simulations (Fig 3C and 3D). The overlap between the eigenvectors increases as the imaginary part of the eigenvalues, β , decreases—equivalent to shrinking the imaginary diameter in the large upper triangular connectivity matrix, $\widetilde{\mathbf{W}}$. The eigenspectrum's real elements, α and γ , have a similar effect: the larger their difference the more the eigenvectors are aligned. Interestingly, when $\alpha = \gamma$, the real distribution has no influence over the overlap between the eigenvectors, such as seen when we decrease the single real-valued distribution (S2(A) Fig, bottom)—the decrease in amplification levels is mostly due to the faster decay times when decreasing the single real-value distribution (S2(A) Fig, top). Because of the square, the larger the absolute value of $\alpha - \gamma$ (negative or positive) the less the eigenvectors are aligned, which we confirm with a large upper triangular matrix, $\widetilde{\mathbf{W}}$, varying the diameter of the eigenspectrum's real distribution while keeping its maximum value at $\max_{\lambda} \text{Re}(\lambda) = 0.5$ (S2(C) Fig). The eigenspectrum's real distribution may affect the level of amplification in two ways: it changes the asymptotic behaviour (decay times) and the eigenvector alignment. These results give a broad intuition for the distinct contribution of the eigenspectrum's imaginary diameter and feedforward norm, but we still do not exactly know how the dynamics of such networks evolve. Thus, we next study the relative directions of the eigenvectors in state-space.

The geometry of output trajectories

If most eigenvectors are pointing in similar directions, the dynamics will be biased towards these directions too (Eq 3). This does not mean that \mathbf{W} or the eigenvector matrix \mathbf{V} are not full rank—on the contrary, they almost always are. What it means is that, in order to quantify the global eigenvector geometry, we have to use the effective rank of \mathbf{V} . The effective rank of \mathbf{V} measures the average number of significant dimensions in its range, and is formally defined as

the exponential of the spectral entropy of its normalised singular values [24]. Specifically, if $\sigma_1, \sigma_2, \dots, \sigma_N$ are the singular values of \mathbf{V} , and $p_i = \frac{\sigma_i}{\|\sigma\|_1}$, with $\|\sigma\|_1 = \sum_{k=1}^N |\sigma_k|$, then

$$\text{erank}(\mathbf{V}) = \exp[H(p_1, \dots, p_N)], \tag{5}$$

where $H(p_1, \dots, p_N)$ is the Shannon entropy, i.e., $H(p_1, \dots, p_N) = -\sum_{k=1}^N p_k \log p_k$.

The effective rank of \mathbf{V} is indeed small in the highly amplifying regimes (Fig 4A), revealing an underlying duality between amplification and output dimensionality. The consequence for the dynamics is that, even though the system may amplify many initial conditions, they nevertheless evolve in the same low dimensional subspace [22]. To identify the dimensionality of this subspace we compute the effective rank of the matrix \mathcal{P} which is constructed as follows: the j^{th} column of \mathcal{P} is the first principal vector of the dynamics, given the j^{th} amplified initial condition of the \mathcal{I}_B basis. We find that there is a discrepancy between the number of amplified directions and effective rank when the system produces large amplifications (Fig 4B; see S3 Fig for results with larger number of principal components). This suggests that the dynamical responses evoked by orthogonal initial conditions evolve in the same subspace, indicating that any added noise will be amplified in the same subspace and that different initial conditions could potentially lead to similar linear readouts (see ref. [6] for capacity estimates of transient amplifying networks). There is thus a trade-off between the number of orthogonal amplified conditions and the noise robustness of the system. The effective rank of \mathbf{V} is preserved in a recurrent network that is a rotated version of an upper triangular one (S4(A) Fig) while the effective rank of the connectivity matrices differ (S4(B) Fig). Similarly, the discrepancy between the number of amplified directions and effective rank is preserved for the rotated (recurrent) weight matrix (S4(C) Fig; compare to Fig 4B), which highlights the robustness of our method to describe the geometry of the network dynamics.

To further describe the system, we use the timescale of the transient amplification, i.e., period, Δt , for which $\|\mathbf{r}(t)\| \geq 1$ for the nonlinear network (Fig 5A and S5 Fig). This timescale

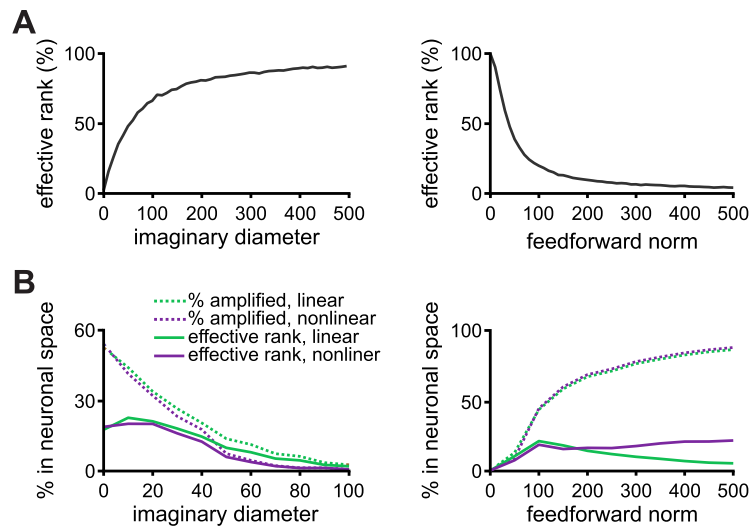


Fig 4. Geometry of output trajectories. **A**, The effective rank of the eigenvector matrix \mathbf{V} of $\tilde{\mathbf{W}}$ as a function of the imaginary diameter (left) and the feedforward norm (right). **B**, Amplified directions and effective rank of the matrix \mathcal{P} (see text) in the linear and nonlinear cases as a function of the imaginary diameter (left) and the feedforward norm (right). The feedforward structure is random from a uniform distribution, and the real distribution is uniform on $(-0.5, 0.5)$. In all cases the network size is $N = 200$. The feedforward Frobenius norm is fixed at 75 for the plots with varying imaginary diameter. The imaginary diameter is fixed at 20 for the plots with varying feedforward norm.

<https://doi.org/10.1371/journal.pcbi.1010365.g004>




Fig 5. Amplification regimes. **A**, Representative examples of non-normal amplification defined by the timescale of the transient response of the nonlinear network—period, Δt , for which $\|\mathbf{r}(t)\| \geq 1$: “weak” ($\Delta t \leq 500$ ms); “short transient” ($500 < \Delta t < 2000$ ms); and “long transient” ($\Delta t \geq 2000$ ms). Grey dotted line indicates $\|\mathbf{r}\| = 1$. **B**, Timescale of the response in the nonlinear network (as in panel A), parametrised by the norms of the spectrum and feedforward structure. Yellow indicates timescale longer than 10 seconds. Boxes correspond to the values used for the plots in panel A (colour coded): (feedforward norm, spectrum norm) = (100, 700), (500, 500), and (700, 100) for weak, short transient, and long transient, respectively. **C**, Maximum norm of the dynamical response per initial condition for different percentages of the norm assigned to the spectrum, ranging from a matrix whose entire norm is assigned to the spectrum (yellow; 100% case, normal matrix) to a matrix whose entire norm is assigned to the feedforward part (dark red; 0% case, nilpotent matrix). The network size is $N = 200$ in all panels. Both eigenspectrum and feedforward structures are random uniform.

<https://doi.org/10.1371/journal.pcbi.1010365.g005>

varies continuously as a function of the norms of the eigenspectrum and the feedforward structure (Fig 5B and S6(A)–S6(F) Fig). Importantly, the dynamics evoked in each of these regimes—defined here by the transient amplification period—lie in different subspaces. For very short periods (e.g., $\Delta t \leq 500$ ms; Fig 5A “weak”), the eigenvectors are effectively orthogonal to each other but span the entire output space equally. For short transient periods (e.g., $500 < \Delta t < 2000$ ms; Fig 5A “short transient”), there is a good balance between amplification of orthogonal inputs and diversity in the responses. For long transient periods (e.g., $\Delta t \geq 2000$ ms; Fig 5A “long transient”), many initial conditions are amplified but the responses lie in the same low-dimensional subspace. This result is well explained by the alignment of the eigenvectors of the simplified 3-by-3 upper triangular matrix (Eq 4): larger feedforward norm or smaller eigenspectrum norm result in eigenvectors being more aligned (S6(G)–S6(I) Fig). Indeed, when we fix the norm of \mathbf{W} , and distribute a—continuously decreasing—percentage of this norm on the diagonal and the rest on the feedforward structure, the network transitions from weakly to strongly amplifying (Fig 5C). Thus, it’s the relation between the diagonal (representing the spectrum) and feedforward parts of the matrix that shapes the dynamics of the network.

Towards biologically realistic networks that satisfy Dale’s law

Up to here we explored different regimes of transient amplification in networks defined by an upper triangular connectivity matrix. This allowed us to have precise control over both eigenspectrum and feedforward structure of the connectivity matrix. However, these abstract networks do not represent biologically realistic neural networks, especially because of Dale’s law, i.e., neurons are either excitatory or inhibitory. In an upper triangular connectivity matrix, $\widetilde{\mathbf{W}}$, the feedforward structure (and consequently feedforward norm) is given by the elements of its strictly upper triangular part, while the eigenspectrum (and consequently spectrum norm) is given by the elements of its diagonal. However, both the eigenspectrum and the feedforward structures are not trivially manipulated in a biological connectivity matrix, \mathbf{W}^B . As a last application, we thus explore how to navigate the regimes of transient amplification in biological networks (i.e., satisfying Dale’s law) based on our results from upper triangular matrices. First, we consider the effect of the absence of self loops in the connectivity matrix. We then focus on

the indirect manipulation of the spectrum norm and the distribution of eigenvalues via modifications of global inhibition and the indirect manipulation of the eigenspectrum's imaginary diameter via modifications of the ratio between excitatory-to-inhibitory and excitatory-to-excitatory connections, respectively.

In the simulations of networks satisfying Dale's law, half of the neurons are chosen to be excitatory (only positive output weights) and the other half to be inhibitory (only negative output weights) [2, 8, 12]. For additional biological plausibility, the connections are sparse, i.e., elements are set to zero following an uniform distribution. We thus have a set of constraints to fulfil: (i) stable asymptotic dynamics ($\lim_{t \rightarrow \infty} \|\mathbf{r}(t)\| = 0$, i.e., $\max_{\lambda} \text{Re}(\lambda) < 1$); (ii) transient amplification ($\|\mathbf{r}(t)\| > \|\mathbf{r}(0)\|$ for $t_0 < t < t_0 + \Delta t$, with $t_0 \geq 0$ and $\Delta t > 0$); and (iii) neurons are excitatory or inhibitory. To do so we build biological weight matrices, \mathbf{W}^{B} , with the algorithm from Stability-Optimised Circuits (SOC) [8] (see [Methods](#)), which fulfils all constraints mentioned above. The SOC algorithm optimises inhibitory connections to ensure that no eigenvalue has real part greater than α_{max} , $\alpha_{\text{max}} < 1$ ($\max_{\lambda} \text{Re}(\lambda) < \alpha_{\text{max}}$). Moreover, it allows for additional constraints to be implemented, such as the level of global inhibition, and whether neurons have autapses, which we use to translate the results from upper triangular matrices, $\widetilde{\mathbf{W}}$, to biological matrices, \mathbf{W}^{B} .

The absence of neuronal self loops shrinks the real distribution. Experimental evidence supports the existence of autapses [25, 26], yet it is common for modelling work to impose no self connections [27, 28]. As a starting point, we thus analysed the consequence of the lack of self connections in networks, i.e., networks without autapses. When neurons are not structurally connected to themselves, the trace of the weight matrix of such a network is equal to zero. This unfolds as follows:

$$\text{Tr}(\mathbf{W}^{\text{B}}) = \sum_{i=1}^N \lambda_i = i \sum_{i=1}^N \text{Im}(\lambda_i) + \sum_{i=1}^N \text{Re}(\lambda_i) = 0. \quad (6)$$

Given that $\sum_{i=1}^N \text{Im}(\lambda_i) = 0$ due the conjugacy of the eigenvalues, the weight matrix \mathbf{W}^{B} without self loops has $\sum_{i=1}^N \text{Re}(\lambda_i) = 0$. This, together with the stability constraint, $\max_{\lambda} \text{Re}(\lambda) = \alpha_{\text{max}} < 1$, bounds the real distribution from below and above, restricting it to a limited diameter that is less than $\alpha_{\text{max}}N$. The maximum range for the real distribution is bounded by $-\alpha_{\text{max}}(N-1)$ and α_{max} , when all real parts but one—defined here as the outlier—are equal to α_{max} and the outlier is equal to $-\alpha_{\text{max}}(N-1)$. This observation explains why the spectrum of a stability-optimised circuit [8], which doesn't have self loops, has an elongated shape along the imaginary axis after optimisation. Not only the positive but also the negative real parts of the eigenspectrum are pushed towards the stability line after optimisation, with the exception of the outlier, which has a large negative value due to the non-self loops constraint and inhibition dominating over excitation [8]. The zero trace condition (particularly $\sum_{i=1}^N \text{Re}(\lambda_i) = 0$) is necessary but not a sufficient condition for the absence of self loops, yet it provides an intuition for its effect on networks defined by upper triangular matrix that can be translated to biological networks.

The existence of the negative outlier together with the zero trace condition has an interesting effect in upper triangular matrices. The larger the value of the outlier (in absolute value), the bigger the amplification ([Fig 6A](#)) and the number of amplified directions ([Fig 6B](#)). On one hand, this can be explained by the fact that more real parts are pushed to the right, creating longer decay envelopes, hence prolonging the time for the hidden feedforward structure to be amplified. On the other hand, this is not the sole source of the increased amplification; the combination of a large negative outlier with the zero trace condition has an additional

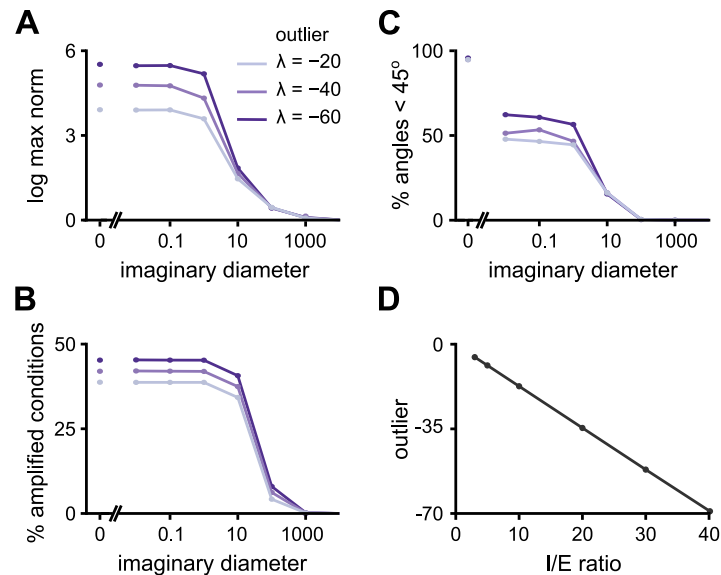


Fig 6. A large negative outlier increases amplification in upper triangular matrices without self loops and is proportional to I/E ratio in networks satisfying Dale's law. A, Maximum response norm for the preferred initial condition as a function of the imaginary diameter using upper triangular connectivity matrices with the zero trace condition and different outliers (coloured coded). The network size is $N = 200$ and the feedforward Frobenius norm is set to 75 in all cases. B, Percentage of directions whose norm is amplified more than 50% as function of the imaginary diameter as in panel A. C, The percentage of angles, between pairs of eigenvectors, that are less than 45° , as a function of the imaginary diameter as in panel A. D, Position of the outlier as a function of the I/E ratio for a network with 100 excitatory and 100 inhibitory neurons sparsely connected with no self loops. An initially random network is optimised with the Stability-Optimised Circuit (SOC) algorithm [8] with $I/E = 40$ (see Methods). The additional outliers are calculated by linearly scaling all inhibitory weights to $I/E = 3, 5, 10, 20, 40$.

<https://doi.org/10.1371/journal.pcbi.1010365.g006>

non-intuitive effect on the geometry of the eigenvectors, i.e., it gives rise to larger eigenvector overlaps (Fig 6C). Without the zero trace condition, increasing the (negative) value of the outlier does not give rise to the same levels of amplification (S7 Fig), confirming that large amplification is an effect from the combination of the negative outlier and the zero trace condition. Interestingly, the negative outlier, γ , reflects the level of global inhibition [2, 8, 12, 29], such that (see Methods)

$$\gamma \approx E \left[1 - \left(\frac{I}{E} \right) \right], \quad (7)$$

where E and I are the sum of excitatory and inhibitory weights per postsynaptic neuron, respectively (Fig 6D).

The eigenspectrum norm is influenced by global inhibition. Taking into account that a large outlier reflects a connectivity matrix with inhibition dominating over excitation [12], i.e., the mean of inhibitory weights is larger than the mean of excitatory weights (Fig 6D), we tested whether strengthening the inhibitory weights would have a similar effect as increasing the (absolute) value of the outlier. For that, we built networks with excitatory and inhibitory neurons with an initial spectral radius $R_{\text{outer}} = 10$ and various ratios of inhibition and excitation, $(\frac{I}{E})$. Non-zero excitatory and inhibitory weights are defined as w_0/\sqrt{N} and $-(\frac{I}{E})w_0/\sqrt{N}$,

respectively, with

$$w_0 = \frac{R_{\text{outer}}}{\sqrt{\frac{p(1-p)}{2} \left[1 + \left(\frac{I}{E} \right) \right]}}, \quad (8)$$

where p is the probability of a connection being non-zero, taken from a uniform distribution. As a result, the eigenspectrum is mostly distributed inside a circle in the complex plane of radius R_{outer} [2] with one negative outlier due to inhibition dominating over excitation [12], given by (see [Methods](#))

$$\gamma \approx \frac{R_{\text{outer}} p \left[1 - \left(\frac{I}{E} \right) \right]}{\sqrt{\frac{2p(1-p)}{N} \left[1 + \left(\frac{I}{E} \right)^2 \right]}}. \quad (9)$$

Importantly, an inhibition dominated network is characterised by a non-uniform distribution of eigenvalues inside the circle of radius R_{outer} , with a denser region near the origin. The denser region is limited by an inner circle [2] with radius given by (see [Methods](#))

$$R_{\text{inner}} \approx \frac{R_{\text{outer}}}{\frac{1}{2} \left[1 + \left(\frac{I}{E} \right)^2 \right]}. \quad (10)$$

Thus, the eigenspectrum of an inhibition dominated network is characterised by the presence of a negative outlier and a denser region of eigenvalues near the origin ([Fig 7A](#)); properties that are preserved after stabilisation of the weight matrix via the SOC algorithm ([Fig 7B](#)). The outlier's value is not drastically changed by the SOC algorithm and is well captured by the analytical expression from [Eq 9](#) ([Fig 7C](#)). Considering only the imaginary distribution, the outer circle shrinks while the inner circle remains mostly constant ([Fig 7D](#)), the latter being well described by [Eq 10](#).

The existence of a denser region of imaginary elements near the origin suggests an effect similar to shrinking the imaginary distribution. This effect may be augmented by the lack of autapses, which moves the real elements towards the stability maximum to accommodate for the large negative outlier that emerges from inhibition dominating over excitation. We find that, indeed, larger global inhibitory strength leads to more amplified conditions and also to slightly larger amplification per condition when the Frobenius norm of the weight matrix is kept fixed ([Fig 7E](#)). By assigning larger values to the inhibitory weights, the spectrum norm decreases and the feedforward norm increases ([Fig 7F](#)). To highlight this finding, we note that when the inhibitory to excitatory ratio is large, $I/E = 40$, the strength of every nonzero excitatory-to-excitatory connection is 0.08, and yet the network is capable of stronger amplification compared to when $I/E = 3$ in which the nonzero excitatory-to-excitatory weights are set to 1.05 (see [Eq 8](#)).

Finally, the new amplified conditions induced by the strongest inhibition do not share their first principal component directions in their dynamical responses ([Fig 7C](#)), i.e., the noise robustness of the system is not compromised in this case [6]. This is possible because we are still in the short transient regime; the long transient regime cannot be reached by solely increasing the global inhibitory strength, i.e., large feedforward norms are always accompanied by small spectrum norms in stability-optimised circuits ([S8 Fig](#)), which restricts the

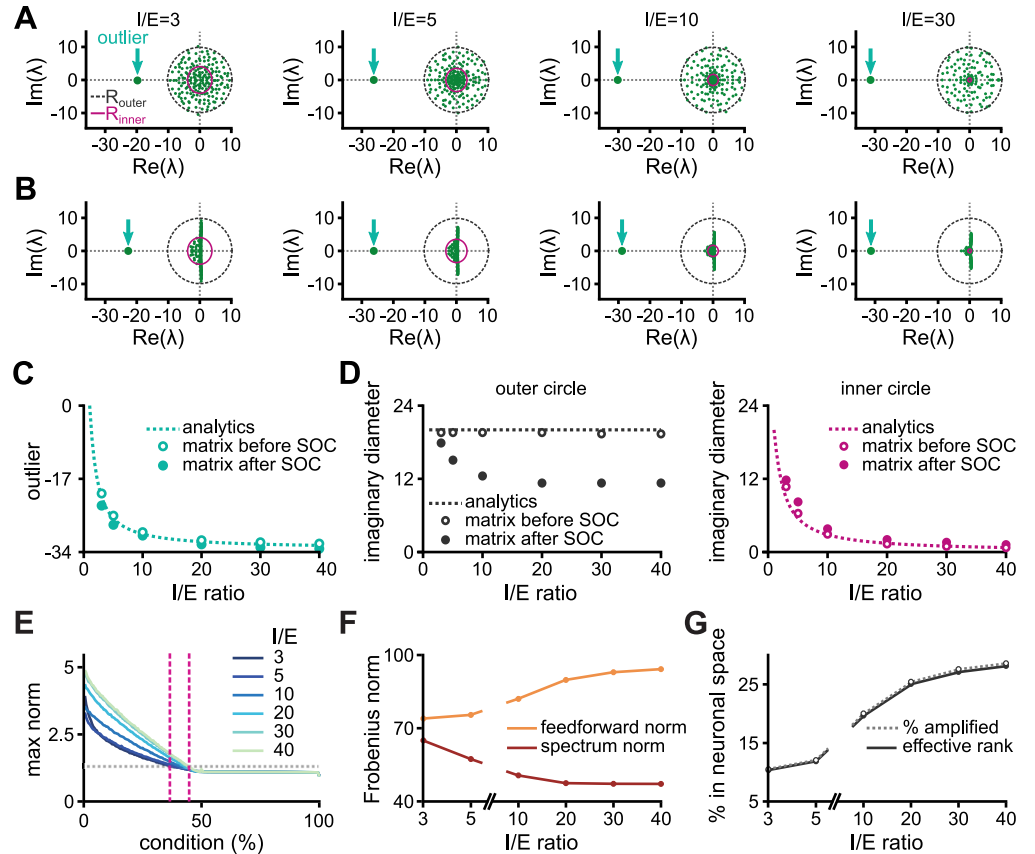


Fig 7. The effect of inhibitory dominance in Dalean matrices. **A**, Eigenspectra of connectivity matrices satisfying Dale’s law: 100 excitatory and 100 inhibitory neurons sparsely connected (probability of connection, $p = 0.1$) without self loops, constructed with spectrum of radius 10 and global inhibitory dominance of strength I/E (indicated on top of each panel). Outlier, outer radius (R_{outer}), and inner circle (R_{inner}) are highlighted. See [Methods](#) for details. **B**, Eigenspectra of connectivity matrices from panel A after optimising inhibitory weights with the SOC algorithm (see [Methods](#) for details). **C**, Value of the purely real outlier before (open circle) and after (closed circle) optimisation. Dashed line represents the analytical expression (Eq 9). **D**, Imaginary diameter of the outer (left) and inner (right) circles. Open and closed circles represent average values before and after SOC optimisation algorithm, respectively, for 1000 random realisations. The outer and inner radii are calculated as the radius for which the density of imaginary elements drops below 0.005 and below half the maximum density, respectively. Dashed grey line (left) indicates $d_{im} = 20$, and purple dashed line (right) represents the analytical expression (Eq 10). **E**, Maximum norm per initial condition for different I/E ratios. Grey dotted line corresponds to a response norm that is 50% larger than the norm of the initial condition. Pink dashed lines indicate the percentage of initial conditions that elicit transients with maximum norm larger than 50% for $I/E = 3$ (lower percentage) and $I/E = 40$ (higher percentage). We linearly scale all weights to keep the same Frobenius norm (equal to 100) for comparison. **F**, The spectrum and the feedforward norms for different values of I/E in the corresponding real Schur transformation. **G**, Percentage of amplified conditions and effective rank of the corresponding matrix \mathcal{P} (defined in the text) in the linear case.

<https://doi.org/10.1371/journal.pcbi.1010365.g007>

accessibility to the different dynamics regimes of transient amplification (S6(D)–S6(F) Fig). Since the overall norm of the matrix stays the same (for comparison reasons), further increasing the inhibitory dominance would unavoidably decrease the excitatory weights even further. Therefore, the amplification power of the network through this mechanism eventually saturates before reaching the long transient regime. Notice that the Frobenius norm of the weight matrix was constrained by our choice of the initial spectral radius, R_{outer} , that, consequently, constrains the possible spectrum and feedforward norms.

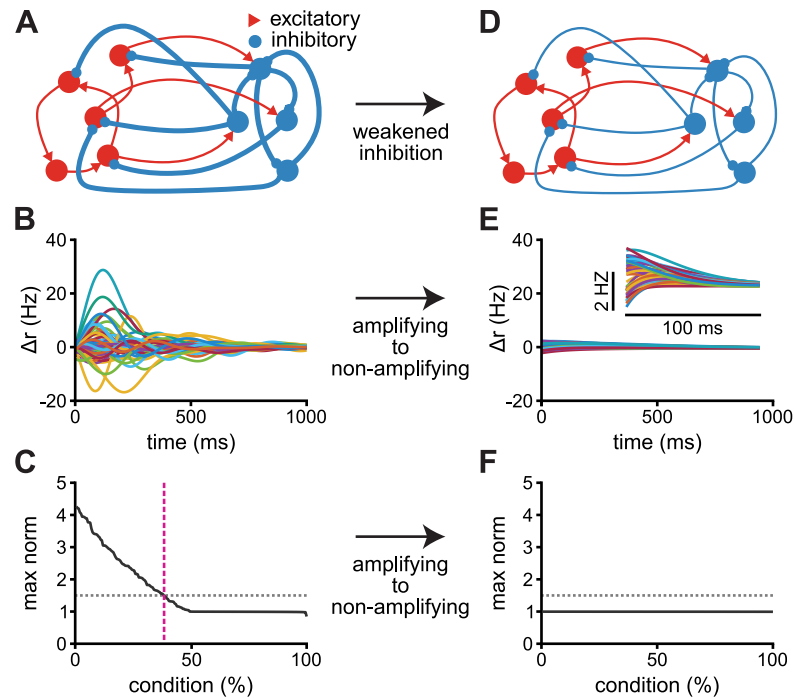


Fig 8. From short transient to non-amplifying with weakened inhibition. **A**, Schematic of a network with strong excitatory and inhibitory connections and I/E ratio of 40. **B**, The network's dynamics given the preferred initialisation. The resulting network is in the short transient regime; the preferred initialisation yields amplifying dynamics. **C**, Maximum response norm for all orthogonal conditions, in decreasing order. Grey dotted line corresponds to a response norm that is 50% larger than the norm of the initial condition. Pink dashed line indicates the percentage of initial conditions with maximum norm larger than 50% of the initial condition. **D**, Schematic of the same network from panel A, but the inhibitory weights are scaled down by a factor of 40 (yielding an I/E ratio of 1), which could be interpreted as the resulting effect of modulation of the inhibitory neurons (or synapses). **E**, The dynamical response given the preferred initialisation; inset depicts the same dynamics on a different scale. **F**, Maximum response norm per condition. The network is unable to amplify any inputs. The maximum norm of the dynamics is equal to the norm of the initial condition (set to be 1) for all initialisations. The network is composed by 100 excitatory and 100 inhibitory neurons sparsely connected and without self loops. The schematics (panels A and D) is adapted from ref. [8].

<https://doi.org/10.1371/journal.pcbi.1010365.g008>

These results indicate that the strength of the global inhibition, if modulated by an external signal, could thus serve as a dynamical switch between amplifying and non-amplifying regimes, and counter-intuitively, weakening of inhibitory synaptic weights would decrease the level of amplification of an originally amplifying network. To test this hypothesis, we built a network with strong inhibitory connections ($I/E = 40$) that enables the network to amplify certain inputs in the short transient regime (Fig 8A–8C). This network transitions into a network that is unable to amplify any inputs by scaling its inhibitory weights down (Fig 8D–8F), confirming that global inhibition can indeed serve as an external switch that controls whether the network is able or not to transiently amplify inputs.

The diameter of the imaginary distribution is influenced by the relationship between excitatory-to-inhibitory and excitatory-to-excitatory connections. In the upper-triangular version of the weight matrix, $\tilde{\mathbf{W}}$, the eigenspectrum's imaginary distribution is given by the off-diagonal terms of $\tilde{\mathbf{W}}$ (see Fig 1B). However, for a biologically plausible weight matrix, \mathbf{W}^B , the manipulation of the imaginary distribution is not as trivial [2, 12]. To get an intuition of the role of each connection type, we use a mean-field approach in which a large weight matrix is simplified as a 2-by-2 matrix [10]: rows and columns correspond to the excitatory and inhibitory populations (Fig 9A), and the elements represent their mean connections (Fig 9B). We

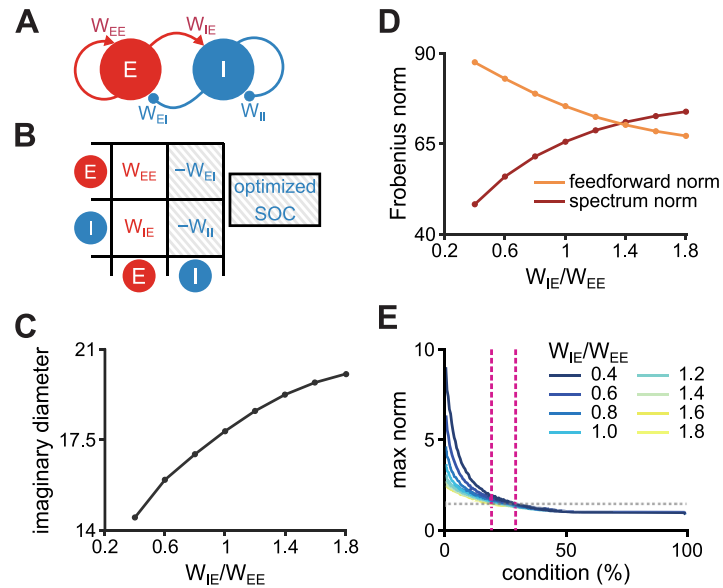


Fig 9. The relationship between E-to-I and E-to-E connectivity strengths alters the imaginary distribution of eigenspectra of Dalean networks. **A**, Schematics of the mean-field analysis of a network with a group of excitatory (E) and a group of inhibitory (I) neurons. The mean weight from E-to-E, E-to-I, I-to-E, and I-to-I are represented by W_{EE} , W_{IE} , W_{EI} , and W_{II} , respectively. **B**, Weight matrix of a simplified network from panel A [10]. Inhibitory connections are optimised by the SOC algorithm [8]. **C**, Imaginary diameter of a network with 100 excitatory and 100 inhibitory neurons as a function of the ratio E-to-I to E-to-E weights for the same networks from panel C. **D**, Feedforward (orange) and spectrum (red) norm as a function of the ratio E-to-I to E-to-E weights for the same networks from panel C. **E**, Maximum norm per initial condition for different ratios E-to-I to E-to-E weights. Grey dotted line indicates response norm that is 50% larger than the norm of the initial condition. Pink dashed lines indicate the maximum percentage of orthogonal initial conditions that evoke response norm 50% larger than initial condition for $W_{IE}/W_{EE} = 1.8$ (lower percentage) and $W_{IE}/W_{EE} = 0.4$ (higher percentage).

<https://doi.org/10.1371/journal.pcbi.1010365.g009>

define W_{EI} , W_{IE} , W_{EE} , and W_{II} (all positive) as the mean connectivity strength of inhibitory-to-excitatory (I-to-E), excitatory-to-inhibitory (E-to-I), excitatory-to-excitatory (E-to-E), and inhibitory-to-inhibitory (I-to-I) groups, respectively. The eigenvalues are complex when the condition $4W_{EI}W_{IE} > (W_{EE} + W_{II})^2$ is satisfied (see Methods for detailed calculation). Real and imaginary parts of the two eigenvalues corresponding to the simplified matrix are thus

$$\text{Re}(\lambda_1) = \text{Re}(\lambda_2) = \frac{1}{2}(W_{EE} - W_{II}) \quad \text{and} \quad (11)$$

$$\text{Im}(\lambda_{1,2}) = \pm \frac{1}{2} \sqrt{4W_{EI}W_{IE} - (W_{EE} + W_{II})^2}. \quad (12)$$

When inhibitory connections, W_{EI} and W_{II} , are optimised for stability [8], we are left with excitatory connections to manipulate the imaginary diameter of the eigenspectrum. Weakening E-to-I, W_{IE} , or strengthening E-to-E, W_{EE} , should shrink the imaginary distribution according to Eq 12.

To test the intuition from the mean-field analysis, we built networks with connectivity matrices obeying Dale’s law and without self loops according to the Stability-Optimised Circuit (SOC) algorithm [8] with varying ratios W_{IE}/W_{EE} . For each network initialisation, we systematically increased the strength of E-to-I connections, W_{IE} , while keeping the same strength of E-to-E connections, W_{EE} . We then optimised inhibitory connections (both W_{EI} and W_{II}) with the SOC algorithm, and linearly scaled all weights to maintain the same Frobenius norm

of the connectivity matrix, $\|\mathbf{W}^B\|$, for all realisations. We found that, in line with Eq 12, the imaginary diameter of the Dalean weight matrix enlarges with the increase of the ratio W_{IE}/W_{EE} (Fig 9C). As a consequence, the spectrum norm increases as the feedforward norm decreases with the strengthening of the E-to-I compared to E-to-E connections (Fig 9D). This indicates that in a stabilised network with connectivity following Dale's law, i.e., when inhibitory connections are optimised for stability, changing the relationship between E-to-I and E-to-E connectivity can influence the diameter of the imaginary distribution of eigenvalues. Due to a decrease in the imaginary diameter, the max norm of dynamics evoked by the orthogonal initial conditions are slightly larger when E-to-I are weaker than E-to-E weights (Fig 9E). These results indicate that excitatory-to-inhibitory connections may play a role in shaping the network's dynamics through their effect on the imaginary distribution of the connectivity matrix.

Discussion

In this article we used upper triangular matrices as abstract representations of the dynamical properties of a connectivity matrix to control the quantities that are relevant for the neural dynamics in the transient amplification regime. By by-passing, temporarily, the connectivity matrix and focusing on a hypothetical Schur transformation, we found new dynamical regimes of large amplification that translated into long transients in non-linear networks. We showed that the amount of transient amplification that a network can produce may be controlled by the ratio between the norms of the spectrum and hidden feedforward structure. Increasing the feedforward norm or decreasing the eigenspectrum's imaginary diameter resulted in larger amplifying dynamics.

Different combinations of the eigenspectrum's and feedforward norms resulted in regimes of transient amplification that lay in subspaces with different dimensionalities and evolved with distinct timescales for non-linear networks. Very short transient periods in non-linear networks spanned the entire output space but only a few orthogonal conditions elicited large amplification. Short transient periods in non-linear networks lay in a lower subspace, but more orthogonal conditions elicited large amplification. Finally, long transient periods in non-linear networks lay in the same low-dimensional subspace, but many orthogonal initial conditions evoked large amplification. The real distribution could considerably change the levels of amplification, but did not affect the relationship between amplification and spectrum or feedforward norms. When the majority of real parts were close to the upper bound of stability, a zero trace condition had an interesting effect: the larger the negative outlier, the closer the real parts were to the stability limit, thus creating larger amplification levels. Moreover, we found a trade-off between the number of orthogonal initial conditions that elicit large amplifications (larger than 50% of initial condition) and the dimensionality of the underlying neuronal dynamics. The source of amplification, i.e., the overlaps of the eigenvectors, inevitably restricted the subspace in which the dynamical outputs evolve, indicating that any noise added to the system is also amplified in the same subspace, giving rise to low robustness to noise [6].

In biologically plausible networks that satisfy Dale's law (i.e., with excitatory and inhibitory neurons), we found that stronger global inhibitory dominance controlled the spectrum norm due to its relationship with the eigenspectrum outlier and a non-uniform distribution of the remaining eigenvalues. Additionally to global inhibition affecting the spectrum norm, we also showed that excitatory-to-inhibitory connections, more specifically its relationship to excitatory-to-excitatory connections, can alter the eigenspectrum's imaginary diameter. We could therefore link the results from upper triangular matrices and biologically plausible networks

via two properties related to specific connectivity patterns: global inhibition and excitatory-to-inhibitory connectivity.

Our work complements previous studies relating manipulations of eigenspectrum and their consequence to neuronal dynamics (e.g., [2, 6, 8, 10, 12, 14, 21, 30]), more specifically on the topic of transient amplification (e.g., [6, 8, 10, 12, 21]). Transient amplification is a phenomenon tightly linked to eigenspectrum and non-normality of a connectivity matrix. For the dynamics to be transient, the network's activity must eventually decay back to baseline, which constrains the eigenspectrum's real part to be less than unity—a stability constraint. For the dynamics to amplify certain inputs, the connectivity matrix must be non-normal [6, 8, 10]. Neuronal networks with excitatory and inhibitory neurons are always non-normal, because of the separation of positive and negative columns [10], but non-normality of a connectivity matrix is not sufficient for a network to transiently amplify inputs [6]. A necessary condition in, e.g., random matrices is that the symmetric (real) part of the connectivity matrix has eigenvalues greater than unity [6]. In networks with sparsely connected excitatory and inhibitory neurons, amplification levels increase with the eigenspectrum radius and the connectivity sparsity, which reflects the global strength of excitatory and inhibitory connections [12]. If the radius is greater than the unity, the network's dynamics are unstable, which can be stabilised by iteratively adjusting inhibitory connections to decrease the spectral abscissa until its values is less than one (and therefore stable) [8]. This algorithm, referred to as Stability-Optimised Circuits (SOC), and used here to generate Dalean matrices, is able to optimise networks to elicit largely amplified transients. Interestingly, after optimisation, the eigenspectrum retains a similar imaginary diameter but the eigenspectrum's real part accumulates near the stability limit apart from the outlier [8]. We confirmed that in Dalean networks, the existence of a large negative outlier imposes that the remaining of the eigenspectrum's real elements be pushed near the spectral abscissa in networks without autapses.

Any structure imposed to a connectivity matrix affects its eigenspectrum in a particular way. Excitatory-inhibitory balance constrains the eigenspectrum to a well defined circle for finite-sized networks, with a non-uniform distribution of eigenvalues inside the circle when the distribution of excitatory and inhibitory weights differ [2]. Cyclic connectivity is reflected by eigenvalues lying on a circle (as oppose to inside for random matrices) [14], and the combination of random matrix and a single feedforward chain with purely imaginary eigenvalues of alternating sign results in an eigenspectrum with two centres at $\pm i$, results in large amplification of oscillatory inputs [21]. Complementary to these previous studies, our work explored simple changes in the eigenspectrum of a connectivity matrix, such as the diameter of imaginary distribution, real negative outlier, and zero trace condition. More importantly, we could directly link these simple changes to properties of Dalean matrices, shining a light on how such changes might be implemented in biologically plausible networks.

Our results in matrices following Dale's law allowed us to find a potential role for changes in global inhibition in neuronal networks. Quick changes in the efficacy of inhibitory synapses, both inhibitory-to-inhibitory and inhibitory-to-excitatory, could act as a switch in the dynamical state of the network, changing from amplifying to non amplifying by reducing all inhibitory efficacies. This switch may be controlled by neuromodulators such as the neuropeptide proctolin [31], acetylcholine [32], dopamine [33], or serotonin [34], as well as stereotypical connectivity motifs [35, 36] that could influence inhibitory activity in a similar way to weakening or strengthening all inhibitory weights. A modulatory increase in the global levels of inhibition could be thus allowing the system to amplify inputs, facilitating signal processing in the brain. This effect resembles the *paradoxical* decrease in inhibitory activity as a result of an increase of external (excitatory) input onto inhibitory neurons [37]. We did not implement any external input, and thus the counter-intuitive (or paradoxical) effect originates from

scaling all inhibitory weights up or down. Intuitively, one would expect that by increasing inhibitory weights, the amplitude of the network activity should decrease, and vice-versa. However, we counter-intuitively found the opposite to be true.

The dynamics of real biological neuronal networks, assessed from large-scale recordings of neural activity, has been shown to typically lie in a low-dimensional space [38], i.e., the number of dimensions necessary to explain the majority of the data variance is much smaller than the number of neurons recorded. The low-dimensional dynamics seen in experiments are well described by, e.g., non-linear rate networks (such as the one used in our work), when connections are optimised for either the experiment's task [7] or recorded neuronal activity [39]. Complementary to large-scale recordings of neuronal dynamics, recent technological improvements have allowed researchers to build connectivity maps of large-scale brain regions [40]. Finding the relationship between connectivity and dynamics is thus crucial to understand the function of brain circuitry. The rank of the weight matrix carries information about the dimensionality of the underlying dynamics generated by the network, but it does not necessarily indicate the dimensionality of the subspace in which the dynamics lie. When, e.g., a weight matrix is constructed as a sum of a unit-rank (generated from two independent vectors) and a random (zero mean) matrix, the underlying dynamics evolves along the direction of the vector used to generate the unit-rank matrix [22]. These vectors are the eigenvectors of the unit-rank matrix [22], and because of the construction of the weight matrix, both weight and eigenvector matrices will share the same effective rank. We have shown a more general case in which the effective rank of the eigenvector matrix can carry information about the dimensionality of the dynamics, even when the weight matrix itself has a high effective rank.

We limited our exploration in upper triangular matrices to the effects of simple changes in the spectrum and feedforward structures, by controlling their range and norm. In networks that satisfy Dale's law, excitatory weights were sparsely connected with the same connection strength for all non-zero entries, changing only the ratio between excitatory-to-inhibitory (E-to-I) and excitatory-to-excitatory (E-to-E) connections. In reality, synapses are constantly changing, following stereotypical synaptic plasticity rules that allow learning of memories [5, 41–43] or additional types of computation such as dynamical switch [44] and input categorisation [45]. Most computational models typically explore learning as a process involving changes in excitatory-to-excitatory (E-to-E) connections accompanied by modifications in inhibitory-to-excitatory (I-to-E) connections for stability [41, 42]. Typically, excitatory-to-inhibitory (E-to-I) connections are randomly initialised and kept fixed throughout the simulations (but see ref. [46]). Our results linking E-to-I connections to the eigenspectrum's imaginary distribution, and consequently, amplification levels, suggest a role for such connection type. Future theoretical and experimental work would be thus necessary to design and measure potential plasticity rules for E-to-I connections that, in synergy with E-to-E [47] and I-to-E plasticity [44], may, e.g., generate new classes of activity patterns.

Our work opens the door for the exploration of new questions related to neuronal dynamics, such as how the structure—besides the norm—of the feedforward part as well as how non-uniform imaginary distributions affect the dynamics of biologically plausible networks.

Methods

Details for upper triangular matrix setup

To construct the upper triangular matrices we take inspiration from the real Schur transformation of real matrices. In this form, the matrix is upper triangular with some 2×2 blocks on the diagonal. These blocks have real entries and their eigenvalues are the complex eigenvalues of the initial matrix (a pair of conjugates). We fix the triangular part that is not involved in the

eigenvalue blocks (more details on this below) and assign to it the norm that we wish by scaling its entries. For the manipulation of the spectrum we construct our distributions as follows. To have the pair of complex eigenvalues $\alpha \pm \beta i$ in the spectrum, we add the block

$$\begin{pmatrix} \alpha & -\beta \\ \beta & \alpha \end{pmatrix}$$

along the diagonal. For real eigenvalues, we add the corresponding real value on the diagonal. The resulting matrix is as follows (example of a 5-by-5 upper triangular matrix with 4 complex eigenvalues and 1 real eigenvalue):

$$\widetilde{\mathbf{W}} = \begin{pmatrix} \alpha_1 & -\beta_1 & \phi_1 & \phi_2 & \phi_3 \\ \beta_1 & \alpha_1 & \phi_4 & \phi_5 & \phi_6 \\ 0 & 0 & \alpha_2 & -\beta_2 & \phi_7 \\ 0 & 0 & \beta_2 & \alpha_2 & \phi_8 \\ 0 & 0 & 0 & 0 & \gamma \end{pmatrix} \quad (13)$$

The eigenvalues are $\lambda_{1,2} = \alpha_1 \pm i\beta_1$, $\lambda_{3,4} = \alpha_2 \pm i\beta_2$, and $\lambda_5 = \gamma$, while the feedforward structure is defined by $\{\phi_i\}$, $i = 1, \dots, 8$. Note that all elements of $\widetilde{\mathbf{W}}$ are purely real. For consistency, we chose to keep the percentage of purely real eigenvalues to 3%, i.e., 6 out of the 200 eigenvalues are purely real, and 194 are complex.

Real distributions. The real distributions we use in this work are the following:

- **A single valued distribution:** all real parts are the same and equal to a fixed value. We use $\text{Re}(\lambda_k) = 0$, $\forall k$ in Fig 3A (light pink) and Fig 3B, as well as S2 Fig (pink) and S6(B) and S6(E) Fig; $\text{Re}(\lambda_k) = -0.5$, $\forall k$ in Fig 3A (dark pink) and S2 Fig (dark pink); and $\text{Re}(\lambda_k) = 0.5$, $\forall k$ in S2(A) Fig (light pink).
- **A distribution with a negative outlier:** in this construction we add a purely real negative outlier, λ_{out} , at a specific point in the complex plane. In Fig 6A–6C and S6(C) Fig, a number $M = \lambda_{\text{out}}/\alpha_{\text{max}} = 2\lambda_{\text{out}}$ of eigenvalues are equal to $\alpha_{\text{max}} = 0.5$ and the rest is equal to zero so that the zero trace condition is satisfied. The value of the outlier is indicated in the figure legends. In S7 Fig the real part of all other eigenvalues is equal to zero, and thus the zero trace condition is not satisfied.
- **A uniform distribution on the interval $(-0.5, 0.5)$:** all real parts, except for the last one, are distributed uniformly between the values -0.5 and 0.5 . As before, because of the zero trace condition, we have to add a small outlier (the last real eigenvalue) to complement for the non-zero sum of the rest of the values. We use this type of distribution in Figs 3A (light green), 4 and 5, as well as S1, S3, S4, S5, S6(A), S6(D) and S6(F) and S9 Figs.
- **A uniform distribution on the interval $(0.5 - d_{\text{re}}, 0.5)$:** all real parts are distributed uniformly between the values $0.5 - d_{\text{re}}$ and 0.5 , creating a uniform distribution with diameter d_{re} and a maximum value of 0.5 . We use this type of distribution in Fig 3A (dark green) with $d_{\text{re}} = 10$ and S2(C) Fig with $d_{\text{re}} = \{0, 0.1, 1, 10, 100\}$.

In all cases, the pairings of the real parts with the corresponding imaginary parts are random—except for forcing the conjugacy of eigenvalues, that is, we make sure that the same real part is paired with conjugate imaginary parts. All simulations are run for 200 realisations

(except where noted), with respect to the randomness of the imaginary distribution, and final quantities are averaged across all realisations for plotting.

Feedforward structures. In the simulations shown in Fig 2, S2(A) and S2(B) Fig (pink), and S6(D)–S6(F) Fig, the feedforward structure of the upper triangular matrix is taken to be equal to the upper triangular part of the Schur transform of a fixed Stability-Optimised Circuit (SOC) [8] with $I/E = 3$ (explained below). The SOC is a matrix known to create strong non-normal amplification and its corresponding feedforward structure is not random, but finely tuned to create amplification. In the simulations shown in Figs 3–6, as well as S1, S2(B) (yellow), S3, S4, S5, S6(A)–S6(C), S7 and S9 Figs, the upper triangular part of the matrix is drawn from a uniform distribution on the interval $(-0.5, 0.5)$ and scaled accordingly to match a specific Frobenius norm.

Imaginary clustering at different points. To understand whether the surprising effect that arises from shrinking the imaginary spectrum is due to the clustering of the eigenvalues, we checked what happens when the imaginary parts of the eigenvalues are not uniformly distributed around zero, but clustered around symmetrically displaced points on the imaginary axis (S1(A) Fig inset). In this case the linear responses exhibit an interesting phenomenon, resembling the beats in acoustics (S1(A) Fig). Because the frequencies are close to each other (due to the clustering), the amplitudes of the different neuronal responses are superimposed when phased, creating a response of very high amplitude (which by our definition would count as amplification). Moreover, the differences in the frequencies create an envelope that is modulating this amplitude over time. The nonlinear responses fail to capture most of the interesting dynamics seen linearly and do not amplify to the same extent (S1(B) Fig). The very high frequency makes it impossible for any potentially amplifying mode to drive the rest of the modes and create a large amplified response. Because of this discrepancy between linear and nonlinear behaviour, we do not consider these regimes as amplifying for the purposes of this manuscript. It is worth noting that similar behaviour to the ± 100 example is seen when clustering the imaginary spectrum at different nonzero values (S1(C) Fig).

Eigenvector overlaps

Recall that the eigenvectors are, in general, complex, in conjugate pairs and that in order to compute the overlap between the eigenvectors we need to consider their inner product. The inner product of two complex vectors is defined as

$$\langle \mathbf{a}, \mathbf{b} \rangle = \sum_i a_i \bar{b}_i \quad (14)$$

and the angle, θ , between two complex vectors is given by

$$\cos(\theta) = \frac{\text{Re}(\langle \mathbf{a}, \mathbf{b} \rangle)}{\|\mathbf{a}\| \|\mathbf{b}\|} \quad (15)$$

Therefore, to compute the angles between the eigenvectors we use Eq 15. In particular, we normalise the eigenvectors to unit norm and compute all pairwise angles. Finally, since $\cos(\pi - \theta) = -\cos(\theta)$, when computing the percentage of small eigenvector overlaps (i.e., less than 45°), we consider as angle the minimum angle between θ and $\pi - \theta$. We would like to note here that non-normality depends on the complex inner product between eigenvectors, and not only its real part. However, we have chosen to compute this more intuitive version of an angle between two complex vectors (which is commonly used in the literature) as a characterisation of the amplification dynamics. We compare these results with an alternative computation of the eigenvector overlap in S9 Fig.

Dimensionality of dynamics—Effective rank of eigenvector matrix

Here we briefly explain the intuition behind the effective rank of the eigenvector matrix \mathbf{V} . This is understood as the number of significant dimensions in the range of a matrix. For example, if the effective rank is equal to κ , then a trajectory evoked by a random initial condition in the range of \mathbf{V} is sufficiently approximated by κ dimensions (see section 3 of ref. [24]). The fact that the effective rank of the eigenvector matrix is small indicates that there are a few prevalent directions in the space spanned by the eigenvectors, which indicates that dynamical trajectories will be biased towards a small subspace of the entire eigenvector space. This is further explored and verified with the computation of the dynamical matrix \mathcal{P} defined below.

Construction of the matrix \mathcal{P} . We construct the matrix \mathcal{P} to understand how correlated the dynamics of the network are, given different initial conditions. This matrix represents the prevalent directions of the dynamics, given different initialisations. This is done as follows: after having identified the optimal orthogonal basis of initial conditions \mathcal{I}_B , we initialise the network at each of the vectors in this basis, one at a time. For each such vector, if the induced dynamics are amplified, i.e., if the norm of the rate vector is at some point in time larger than 1.5 (the initialisation vectors have always unit norm), then we perform Principal Component Analysis on the dynamics. More specifically, we compute the eigenvectors of the covariance matrix of the neuronal dynamics for each of these simulations. We only consider the eigenvector corresponding to the largest eigenvalue and store it as a column in the matrix \mathcal{P} . Once we have initialised the network at all vectors in \mathcal{I}_B we are left with a $N \times M$ matrix \mathcal{P} . The number M is the same as the number of conditions that lead to an amplified response and provides a maximum bound for the effective rank of matrix \mathcal{P} .

The effective rank of \mathcal{P} thus gives us the effective dimensionality of the space spanned by the columns of \mathcal{P} . If the effective rank is less than the number of columns, we can deduce that orthogonal initial conditions have first principal components that are closely aligned to each other in state-space. This implies that the initial network amplifies orthogonal initial conditions along the same low dimensional subspace.

In S3 Fig we also compare the effective rank of the matrix \mathcal{P} when the number of principal components stored as columns (for each amplified initialisation) is such that the total variance captured is greater than 85%. In that case, the matrix \mathcal{P} has size $N \times \mathcal{K}$. Here, $\mathcal{K} = \sum_{i=1}^M \kappa_i$, where κ_i is the number of principal components needed to explain at least 85% of the variance of the neuronal response, when initialised at the i^{th} (amplified) condition of \mathcal{I}_B . We find similar results, i.e., in the long transient regime there is a big discrepancy between the total number of columns (\mathcal{K}) and the effective rank of the dynamical matrix \mathcal{P} . It is worth mentioning that the effective rank is bounded by the number of neurons N . The fact that even though the number of columns is much larger than N in the long transient, but the effective rank still fails to reach its bound, verifies the intuition obtained by the main results in Fig 4.

Construction of recurrent networks satisfying Dale's law

The recurrent networks satisfying Dale's law are constructed following the Stability-Optimised Circuit (SOC) algorithm [8]. In our simulations, 50% of the neurons are excitatory and 50% inhibitory, so that the first half columns of \mathbf{W}^B are strictly positive and the second half strictly negative. The connections are initially generated at random from a uniform distribution with probability of connection $p = 0.1$ and individual weights given by: $w_E = w_0/\sqrt{N}$ for excitatory and $w_I = -(\frac{I}{E})w_0/\sqrt{N}$ for inhibitory, where $(\frac{I}{E})$ is the ratio between inhibition and excitation,

and N is the size of the network. The initial weight, w_0 , is defined as

$$w_0 = \frac{R_{\text{outer}}}{\sqrt{\frac{p(1-p)}{2} \left[1 + \left(\frac{I}{E} \right)^2 \right]}}, \quad (16)$$

where $R_{\text{outer}} = 10$ is the radius of the eigenspectrum distribution. After the initialisation of a matrix, we implement the SOC optimisation algorithm, which modifies only inhibitory weights to enforce that the spectral abscissa is less than 0.5 ($\max_{\lambda} \text{Re}(\lambda) = \alpha_{\text{max}} = 0.5$). In this process, both zero and non-zero inhibitory connections are modified. Thus, to maintain a certain level of sparsity in inhibitory connections, we keep the density of inhibitory connections lower than 0.4. Additionally, we keep the absence of self loops and the same I/E ratio by linearly scaling all non-zero inhibitory weights. We impose different conditions to connections depending on which aspect of the connectivity we explore, as define below.

Varying I/E ratios. In Fig 7 and S8 Fig we pre-define the ratio I/E and then optimise the inhibitory connections. In Figs 6D and 8 we optimise the connectivity for $I/E = 40$. We linearly scale all inhibitory weights and calculate the resulting outlier in Fig 6D for $I/E = 30, 20, 10, 5, 3$. We linearly scale all inhibitory weights by the fraction $1/40$ to get a resulting weight matrix with $I/E = 1$ in Fig 8.

Varying the E-to-I and E-to-E connections. In Fig 9 we implement the SOC algorithm with different combinations of the E-to-I and E-to-E connectivity strengths. We define our “standard” network, i.e., with the same E-to-I and E-to-E connectivity strengths ($W_{\text{IE}}/W_{\text{EE}} = 1$) as described above, optimised using $I/E = 3$. For each ratio E-to-I to E-to-E, $W_{\text{IE}}/W_{\text{EE}} = \{0.4, 0.6, 0.8, 1.2, 1.4, 1.6, 1.8\}$, we start with a new matrix with weights defined as above and then linearly scale all E-to-I non-zero weights by $W_{\text{IE}}/W_{\text{EE}}$. After scaling E-to-I weights, we run the SOC algorithm, as described above, keeping $I/E = 3$. Finally, after optimisation we linearly scale all weights so that the Frobenius norm of the resulting \mathbf{W}^{B} is the same as the standard weight matrix, which has $W_{\text{IE}}/W_{\text{EE}} = 1$.

Analysis of a 3-by-3 upper triangular weight matrix

To get an intuition for the role of the eigenspectrum, feedforward norm, and outlier, we analyse a simplified 3-by-3 upper triangular matrix,

$$\widetilde{\mathbf{W}} = \begin{pmatrix} \alpha & \beta & \phi_2 \\ -\beta & \alpha & \phi_1 \\ 0 & 0 & \gamma \end{pmatrix}. \quad (17)$$

This matrix has three eigenvalues, λ_k , given by

$$\begin{cases} \lambda_1 = \gamma \\ \lambda_2 = \alpha - i\beta \\ \lambda_3 = \alpha + i\beta, \end{cases} \quad (18)$$

where α , β , and γ can be interpreted as the average of the real part of the eigenspectrum distribution, the eigenspectrum’s imaginary diameter, and the outlier. The other two non-zero entries, ϕ_1 and ϕ_2 , represent the feedforward structure with norm $\phi_{\text{norm}} = \sqrt{\phi_1^2 + \phi_2^2}$.

The eigenvectors associated with λ_k are

$$\mathbf{v}_1 = \begin{pmatrix} a_1 \\ b_1 \\ c_1 \end{pmatrix} = \begin{pmatrix} \frac{-\alpha\phi_2 + \gamma\phi_2 + \beta\phi_1}{(\alpha - \gamma)^2 + \beta^2} \\ \frac{-\alpha\phi_1 + \gamma\phi_1 - \beta\phi_2}{(\alpha - \gamma)^2 + \beta^2} \\ 1 \end{pmatrix} c_1, \tag{19}$$

$$\mathbf{v}_2 = \begin{pmatrix} a_2 \\ b_2 \\ c_2 \end{pmatrix} = \begin{pmatrix} i \\ 1 \\ 0 \end{pmatrix} b_2, \tag{20}$$

and

$$\mathbf{v}_3 = \begin{pmatrix} a_3 \\ b_3 \\ c_3 \end{pmatrix} = \begin{pmatrix} -i \\ 1 \\ 0 \end{pmatrix} b_3. \tag{21}$$

We can thus calculate the normalised inner product of the vectors,

$$\frac{\langle \mathbf{v}_2, \mathbf{v}_3 \rangle}{\|\mathbf{v}_2\| \|\mathbf{v}_3\|} = 0 \tag{22}$$

and

$$\begin{aligned} \eta &\equiv \frac{\langle \mathbf{v}_1, \mathbf{v}_2 \rangle}{\|\mathbf{v}_1\| \|\mathbf{v}_2\|} = \frac{\langle \mathbf{v}_1, \mathbf{v}_3 \rangle}{\|\mathbf{v}_1\| \|\mathbf{v}_3\|} \\ &= \frac{-[\phi_1(\alpha - \gamma) + \beta\phi_2] + i[\phi_2(\alpha - \gamma) - \beta\phi_1]}{\sqrt{2\{[\phi_2(\alpha - \gamma) - \beta\phi_1]^2 + [\phi_1(\alpha - \gamma) + \beta\phi_2]^2 + [(\alpha - \gamma)^2 + \beta^2]^2\}}} \\ &= \frac{-[\phi_1(\alpha - \gamma) + \beta\phi_2] + i[\phi_2(\alpha - \gamma) - \beta\phi_1]}{\sqrt{2[\phi_1^2 + \phi_2^2 + (\alpha - \gamma)^2 + \beta^2][(\alpha - \gamma)^2 + \beta^2]}}. \end{aligned} \tag{23}$$

From Eq 23 we can calculate the absolute value of the normalised inner product of the non-orthogonal eigenvectors (Eq 4 in the main text),

$$|\eta| = \sqrt{\eta\bar{\eta}} = \frac{\phi_{\text{norm}}}{\sqrt{2[\phi_{\text{norm}}^2 + (\alpha - \gamma)^2 + \beta^2]}}. \tag{24}$$

We plot $|\eta|$ (from Eq 24) as a function of β in Fig 3C and as a function of ϕ_{norm} in Fig 3D (values of α and γ are described in the caption). We also plot $|\eta|$ (Eq 24) as a function of β and ϕ_{norm} representing the spectrum’s imaginary diameter and feedforward norm, respectively, in S6(G)–S6(I) Fig.

Mean-field analysis of a 2-by-2 weight matrix with excitatory and inhibitory neuronal populations

To get an intuition for the role of the real and imaginary distributions of a Dalean weight matrix, we analyse a 2-by-2 matrix

$$\mathbf{W}^B = \begin{pmatrix} W_{EE} & -W_{EI} \\ W_{IE} & -W_{II} \end{pmatrix}, \tag{25}$$

where W_{EE} , W_{EI} , W_{IE} , and W_{II} correspond to the mean excitatory-to-excitatory (E-to-E), inhibitory-to-excitatory (I-to-E), excitatory-to-inhibitory (E-to-I), and inhibitory-to-inhibitory (I-to-I) connections, respectively. The eigenvalues of this connectivity matrix are

$$\begin{cases} \lambda_1 = \frac{1}{2} \left[W_{EE} - W_{II} + \sqrt{(W_{EE} + W_{II})^2 - 4W_{EI}W_{IE}} \right] \\ \lambda_2 = \frac{1}{2} \left[W_{EE} - W_{II} - \sqrt{(W_{EE} + W_{II})^2 - 4W_{EI}W_{IE}} \right]. \end{cases} \tag{26}$$

The eigenvalues are complex when

$$4W_{EI}W_{IE} > (W_{EE} + W_{II})^2, \tag{27}$$

with real part $\text{Re}(\lambda_1) = \text{Re}(\lambda_2) = \frac{1}{2}(W_{EE} - W_{II})$ and imaginary part $\text{Im}(\lambda_1) = \frac{1}{2}\sqrt{4W_{EI}W_{IE} - (W_{EE} + W_{II})^2}$ and $\text{Im}(\lambda_2) = -\frac{1}{2}\sqrt{4W_{EI}W_{IE} - (W_{EE} + W_{II})^2}$. The imaginary part of both eigenvalues is small when the two quantities from Eq 27 are similar,

$$4W_{EI}W_{IE} \gtrsim (W_{EE} + W_{II})^2, \tag{28}$$

and large when the inter-group connections are much greater than the intra-group connections,

$$4W_{EI}W_{IE} \gg (W_{EE} + W_{II})^2. \tag{29}$$

Therefore, in realistic Dalean networks with excitatory and inhibitory neurons, the imaginary diameter can be controlled by the inter-group connectivity strength, with small imaginary diameters by enforcing the relationship between mean weights according to Eq 28.

Relationship between negative outlier and inhibition-to-excitation ratio

To find the relationship between the negative outlier and the inhibition-to-excitation ratio, I/E , we consider a special case in which the sum of all rows of a biological weight matrix, \mathbf{W}^B , are equal to γ ,

$$\sum_{j=1}^N w_{ij}^B = \gamma, \forall i, \tag{30}$$

where N is the number of neurons. Additionally, we consider that the sum of excitatory and inhibitory weights in each row is the same,

$$\sum_{j=1}^{N/2} w_{ij}^B = E, \sum_{j=\frac{N}{2}+1}^N w_{ij}^B = -I, \text{ and } E - I = \gamma, \forall i. \tag{31}$$

In this particular case the outlier is

$$\gamma = E - I = E \left(1 - \frac{I}{E} \right). \tag{32}$$

We extrapolate this to a random weight matrix with structured inhibition as

$$\gamma \propto 1 - \left(\frac{I}{E} \right), \tag{33}$$

which is confirmed in Fig 6D.

For Dalean matrices we used the SOC algorithm, with non-negative excitatory weights equal to w_0/\sqrt{N} and inhibitory weights equal to $-(\frac{I}{E})w_0/\sqrt{N}$. Non-negative connections were randomly assigned with probability p from a uniform distribution, and thus the total excitatory (E) and inhibitory (I) input weights per neuron are

$$E_i = \sum_{j=1}^{N/2} w_{ij}^B \approx E = \frac{Npw_0}{2\sqrt{N}}, \text{ and } I_i = - \sum_{j=\frac{N}{2}+1}^N w_{ij}^B \approx I = - \left(\frac{I}{E} \right) \frac{Npw_0}{2\sqrt{N}}. \tag{34}$$

The initial weight, w_0 , is based on the eigenspectrum’s radius, R_{outer} (Eq 16), which results in the outlier being given by

$$\gamma \approx \frac{pw_0\sqrt{N}}{2} \left[1 - \left(\frac{I}{E} \right) \right] = \frac{R_{\text{outer}}p \left[1 - \left(\frac{I}{E} \right) \right]}{\sqrt{\frac{2p(1-p)}{N} \left[1 + \left(\frac{I}{E} \right)^2 \right]}}. \tag{35}$$

Relationship between the eigenspectrum and inhibition-to-excitation ratio

In the networks with excitatory and inhibitory neurons, we enforce that inhibition dominates over excitation. This means that the variance of the weight distributions are different. In this case, the eigenvalues lie inside the circle of radius [2]

$$R_{\text{outer}} \approx \sqrt{\frac{N}{2} (\sigma_E^2 + \sigma_I^2)} = w_0 \sqrt{\frac{p(1-p)}{2} \left[1 + \left(\frac{I}{E} \right)^2 \right]}, \tag{36}$$

where σ_E^2 and σ_I^2 are the variance of the excitatory and inhibitory weight distributions, respectively. Non-zero excitatory and inhibitory weights are equal to w_0/\sqrt{N} and $-(\frac{I}{E})w_0/\sqrt{N}$, respectively, and randomly chosen with probability p drawn from a uniform distribution.

The distribution of eigenvalues is not uniform when the variance of excitatory and inhibitory weight distributions are different [2]. In this case there is an accumulation of eigenvalues inside an inner circle of radius given by $R_{\text{inner}} \approx \sqrt{N}\sigma_{\text{min}}$, where $\sigma_{\text{min}} = \min(\sigma_E, \sigma_I)$. When inhibition is stronger than excitation, the distribution with smaller variance is the excitatory one,

and thus $\sigma_{\min}^2 = \sigma_E^2 = w_0^2 p(1-p)/N$, resulting in an inner radius given by

$$R_{\text{inner}} \approx \sqrt{N} \sigma_{\min} = w_0 \sqrt{p(1-p)} = \frac{R_{\text{outer}}}{\sqrt{\frac{1}{2} \left[1 + \left(\frac{I}{E} \right)^2 \right]}}. \quad (37)$$

Supporting information

S1 Text. Supplemental text. Sections “Why upper triangular?”, “Alternative feedforward structures”, and “Biologically plausible network dynamics with strictly positive rates”. (PDF)

S1 Fig. Imaginary clustering at different points. Dynamical responses of a spectrum that is clustered around the points 100 and -100 with respect to the imaginary axis. The imaginary radius around these points is 0.5. The real distribution is uniform on the interval $(-0.5, 0.5)$. **A**, Linear dynamical response; the network shows an amplified response, effectively due to superposition of almost identical frequencies. Inset: eigenspectrum distribution. **B**, Nonlinear responses are not amplifying; the very high frequency together with the saturation point prevents the network’s modes from driving each other in order to create an amplifying response. **C**, Clustering at other finite points, $\{\pm 10, \pm 30, \pm 50, \pm 70, \pm 90\}$, shows the same discrepancy between the linear and nonlinear behaviour, as measured by the percentage of amplified responses. (EPS)

S2 Fig. Varying the real and feedforward distributions. **A**, Exploring the single valued real distribution. We compare results for three real values: $-0.5, 0$ and 0.5 . Top: maximum response norm for preferred initial condition. Naturally a larger real part leads to more amplification as the decay envelope becomes slower. Middle: % of amplified conditions that are amplified by at least 50%; this is also affected by the value of the real part, indicating that the amplification landscape changes its shape in a uniform way. Bottom: the percentage of pairwise eigenvector angles is independent of the real value, i.e., the increased amount of amplification is mainly a result of the slower decay times. Results in all cases are qualitatively similar in their dependence on the imaginary radius. **B**, Comparing results for two different feed-forward structures. One is the feedforward structure taken from the corresponding feedforward part of a matrix constructed using the Stability-Optimised algorithm [8] (pink). The other has a uniform feedforward entry distribution, with overall feedforward norm equal to the stability-optimised one (yellow). In both cases the spectra are identical and correspond to the spectral distribution of the pink curve from **A**, i.e., single real value at zero, varying imaginary range represented on the x-axis. **C**, Same as panel **A** but the real part of the eigenspectrum is uniformly distributed between $(0.5 - d_{re}, 0.5)$. Values of d_{re} are indicated in the figure (colour coded). (EPS)

S3 Fig. Effective rank of matrix \mathcal{P} for a larger number of principal components. **A**, Comparison of the number of columns and effective rank of the dynamical matrix \mathcal{P} as a function of the imaginary diameter. The feedforward structure is random from a uniform distribution, and scaled to have Frobenius norm equal to 75. Number of neurons, $N = 200$. **B**, Number of columns and effective rank of \mathcal{P} as a function of the feedforward norm. The imaginary diameter is fixed and equal to 20. Number of neurons, $N = 200$. The matrix \mathcal{P} is constructed such that for each amplified initial condition \mathbf{a}_i in \mathcal{I}_B , there are κ_i principal components stored as columns in the matrix \mathcal{P} . The number κ_i is the number of principal components that are

needed to capture at least 85% of the total variance of the dynamical response given condition \mathbf{a}_i . The large discrepancy between the number of columns and the effective rank in the long transient regimes verifies the results shown in the manuscript.

(EPS)

S4 Fig. Comparison between upper triangular and recurrent matrices. **A**, Effective rank of the eigenvector matrix, \mathbf{V} , for an upper triangular matrix (dashed cyan line) and a recurrent matrix (pink solid line) as a function of the imaginary diameter (left) and feedforward norm (right). The recurrent weight matrix, \mathbf{W}_{rec} , is the upper triangular matrix, $\widetilde{\mathbf{W}}$, rotated with an orthonormal basis \mathbf{U} : $\mathbf{W}_{\text{rec}} = \widetilde{\mathbf{W}}\mathbf{U}^\dagger$. **B**, Same as panel A, but the effective rank is calculated for the weight matrix. **C**, Amplified directions and effective rank of the matrix \mathcal{P} in the linear and nonlinear cases for the recurrent matrix in panels A and B as a function of the imaginary diameter (left) and the feedforward norm (right). The feedforward structure is random from a uniform distribution, and the real distribution is uniform on $(-0.5, 0.5)$. In all cases the network size is $N = 200$. The feedforward Frobenius norm is fixed at 75 for the plot varying the imaginary diameter. The imaginary diameter is fixed at 20 for the plot with varying feedforward norm.

(EPS)

S5 Fig. Timescale of the transient dynamics as a function of the eigenspectrum's imaginary diameter (left) and the feedforward norm (right) of an upper triangular matrix. Timescale is defined as the period, Δt , for which $\|\mathbf{r}(t)\| \geq 1$ for the nonlinear network. The feedforward structure is random from a uniform distribution, and the real distribution is uniform on $(-0.5, 0.5)$. The network size is $N = 200$, and the feedforward Frobenius norm is fixed at 75 for the plot varying the eigenspectrum's imaginary diameter (left). The imaginary diameter is fixed at 20 for the plot with varying feedforward norm (right).

(EPS)

S6 Fig. Regimes of transient amplification for distinct eigenspectra and feedforward structures from simulations (A-F) and analytics (G-I). **A-F**, Timescale of response (time for which $\|\mathbf{r}(t)\| \geq 1$) for distinct eigenspectra and feedforward structures (specified below) from simulations with $N = 200$ neurons. **A**, Random uniform eigenspectrum distribution and random uniform feedforward distribution. Same plot as in Fig 5B. **B**, Eigenspectrum's real part fixed at zero, random uniform distribution of the eigenspectrum's imaginary part, and random uniform feedforward distribution. **C**, Inhibitory dominance in the eigenspectrum and random uniform feedforward distribution. **D**, Random uniform eigenspectrum distribution and feedforward distribution from stability-optimised circuit (SOC). **E**, Eigenspectrum's real part fixed at zero, random uniform distribution of the eigenspectrum's imaginary part, and feedforward distribution from SOC. **F**, Purely real eigenspectrum with inhibitory dominance and feedforward distribution from SOC. **G-I**, Normalised inner product of two eigenvectors from a 3-by-3 upper triangular matrix (Eq 4) as a function of feedforward norm and imaginary diameter for $\alpha - \gamma = 0$ (**G**), $\alpha - \gamma = 128$ (**H**), and $\alpha - \gamma = 1024$ (**I**).

(EPS)

S7 Fig. Effect of varying the outlier in the upper triangular setting without the zero trace condition. Maximum response norm for the preferred initial condition (left), percentage of directions whose norm is amplified more than 50% (middle), and percentage of angles, between pairs of eigenvectors, that are less than 45° (right) as a function of the imaginary diameter. In all cases the network size is $N = 200$ and the Frobenius norm of all matrices is normalised to 100. Different colours correspond to four different outlier values (colour coded).

The rest of the eigenspectrum's real distribution is the same: uniformly distributed between -0.5 and 0.5 .

(EPS)

S8 Fig. Feedforward (left) and spectrum (right) norms as a function of excitatory weights and I/E ratio in Dalean matrices without self loops. All matrices were generated with the SOC algorithm [8]. The y-axis corresponds to the value of individual excitatory weights, which are sparsely connected with probability 0.1 (see [Methods](#)). The x-axis corresponds to the I/E ratio, set to be constant by the SOC algorithm. In all cases the network size is $N = 200$ (100 excitatory and 100 inhibitory neurons), and the Frobenius norm of all matrices is normalised to 100.

(EPS)

S9 Fig. Eigenvector overlap comparison. We compare the eigenvector overlap statistics as a function of the imaginary diameter. On one hand we compute the percentage of pairs that form angles $< 45^\circ$. This is defined using the real part of the inner product, i.e., $\cos(\theta) = \text{Re}(\langle \alpha, \beta \rangle)$, for the complex eigenvectors α, β (green). Alternatively, we can compute the percentage of pairs for which the magnitude of the complex-valued inner product $|\langle \alpha, \beta \rangle|$ is larger than 0.7 (pink). Both yield similar results qualitatively.

(EPS)

S10 Fig. Comparing different feedforward structures. Maximum response norm for the preferred initial condition (left), percentage of directions whose norm is amplified more than 50% (middle), and effective rank of the eigenvector matrix (right), as a function of the imaginary diameter. Three feedforward structures are compared. Green: uniform distribution as in [Fig 3A](#) of manuscript. Purple: a feedforward structure limited to chains of length 2, i.e., each Schur unit only connects directly to the next unit. Yellow: a sparse feedforward structure with probability of connection equal to 0.1. For normalisation reasons, the Frobenius norm of all feedforward structures is set to be equal to 75. Moreover, the real distributions of the sparse and 2-chain matrices do not satisfy the zero-trace condition. Inset in panel A shows the uniform and sparse networks again on a more appropriate scale for their values.

(EPS)

Acknowledgments

We thank Friedemann Zenke for his comments, especially on the effect of the self loops on the spectrum. We also thank Ken Miller and Bill Podlaski for helpful comments.

Author Contributions

Conceptualization: Georgia Christodoulou, Tim P. Vogels, Everton J. Agnes.

Formal analysis: Georgia Christodoulou, Everton J. Agnes.

Funding acquisition: Tim P. Vogels.

Investigation: Georgia Christodoulou.

Methodology: Georgia Christodoulou, Tim P. Vogels, Everton J. Agnes.

Project administration: Tim P. Vogels, Everton J. Agnes.

Resources: Tim P. Vogels.

Software: Georgia Christodoulou.

Supervision: Tim P. Vogels, Everton J. Agnes.

Validation: Everton J. Agnes.

Visualization: Georgia Christodoulou, Everton J. Agnes.

Writing – original draft: Georgia Christodoulou, Tim P. Vogels, Everton J. Agnes.

Writing – review & editing: Georgia Christodoulou, Tim P. Vogels, Everton J. Agnes.

References

1. Vogels TP, Rajan K, Abbott LF. Neural network dynamics. *Annual Review of Neuroscience*. 2005; 28:357–376. <https://doi.org/10.1146/annurev.neuro.28.061604.135637> PMID: 16022600
2. Rajan K, Abbott LF. Eigenvalue spectra of random matrices for neural networks. *Physical Review Letters*. 2006; 97:188104. <https://doi.org/10.1103/PhysRevLett.97.188104> PMID: 17155583
3. Sussillo D, Abbott LF. Generating coherent patterns of activity from chaotic neural networks. *Neuron*. 2009; 63:544–557. <https://doi.org/10.1016/j.neuron.2009.07.018> PMID: 19709635
4. Stroud JP, Porter MA, Hennequin G, Vogels TP. Motor primitives in space and time via targeted gain modulation in cortical networks. *Nature Neuroscience*. 2018; 21:1774–1783. <https://doi.org/10.1038/s41593-018-0276-0> PMID: 30482949
5. Susman L, Brenner N, Barak O. Stable memory with unstable synapses. *Nature Communications*. 2019; 10:4441. <https://doi.org/10.1038/s41467-019-12306-2> PMID: 31570719
6. Bondanelli G, Ostojic S. Coding with transient trajectories in recurrent neural networks. *PLOS Computational Biology*. 2020; 16:e1007655. <https://doi.org/10.1371/journal.pcbi.1007655> PMID: 32053594
7. Mante V, Sussillo D, Shenoy KV, Newsome WT. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*. 2013; 503:78–84. <https://doi.org/10.1038/nature12742> PMID: 24201281
8. Hennequin G, Vogels TP, Gerstner W. Optimal control of transient dynamics in balanced networks supports generation of complex movements. *Neuron*. 2014; 82:1394–1406. <https://doi.org/10.1016/j.neuron.2014.04.045> PMID: 24945778
9. Remington ED, Narain D, Hosseini EA, Jazayeri M. Flexible sensorimotor computations through rapid reconfiguration of cortical dynamics. *Neuron*. 2018; 98:1005–1019. <https://doi.org/10.1016/j.neuron.2018.05.020> PMID: 29879384
10. Murphy BK, Miller KD. Balanced amplification: a new mechanism of selective amplification of neural activity patterns. *Neuron*. 2009; 61:635–648. <https://doi.org/10.1016/j.neuron.2009.02.005> PMID: 19249282
11. Chaudhuri R, Bernacchia A, Wang XJ. A diversity of localized timescales in network activity. *eLife*. 2014; 3:e01239. <https://doi.org/10.7554/eLife.01239> PMID: 24448407
12. Hennequin G, Vogels TP, Gerstner W. Non-normal amplification in random balanced neuronal networks. *Physical Review E*. 2012; 86:011909. <https://doi.org/10.1103/PhysRevE.86.011909> PMID: 23005454
13. Elsner L, Paardekooper MHC. On measures of nonnormality of matrices. *Linear Algebra and its Applications*. 1987; 92:107–123. [https://doi.org/10.1016/0024-3795\(87\)90253-9](https://doi.org/10.1016/0024-3795(87)90253-9)
14. Goldman MS. Memory without feedback in a neural network. *Neuron*. 2009; 61:621–634. <https://doi.org/10.1016/j.neuron.2008.12.012> PMID: 19249281
15. Churchland MM, Cunningham JP, Kaufman MT, Ryu SI, Shenoy KV. Cortical preparatory activity: representation of movement or first cog in a dynamical machine? *Neuron*. 2010; 68:387–400. <https://doi.org/10.1016/j.neuron.2010.09.015> PMID: 21040842
16. Churchland MM, Cunningham JP, Kaufman MT, Foster JD, Nuyujukian P, Ryu SI, et al. Neural population dynamics during reaching. *Nature*. 2012; 487:51–56. <https://doi.org/10.1038/nature11129> PMID: 22722855
17. Ames KC, Ryu SI, Shenoy KV. Neural dynamics of reaching following incorrect or absent motor preparation. *Neuron*. 2014; 81:438–451. <https://doi.org/10.1016/j.neuron.2013.11.003> PMID: 24462104
18. Shafi M, Zhou Y, Quintana J, Chow C, Fuster J, Bodner M. Variability in neuronal activity in primate cortex during working memory tasks. *Neuroscience*. 2007; 146:1082–1108. <https://doi.org/10.1016/j.neuroscience.2006.12.072> PMID: 17418956

19. Barak O, Tsodyks M, Romo R. Neuronal population coding of parametric working memory. *Journal of Neuroscience*. 2010; 30:9424–9430. <https://doi.org/10.1523/JNEUROSCI.1875-10.2010> PMID: 20631171
20. Hussar CR, Pasternak T. Memory-guided sensory comparisons in the prefrontal cortex: contribution of putative pyramidal cells and interneurons. *Journal of Neuroscience*. 2012; 32:2747–2761. <https://doi.org/10.1523/JNEUROSCI.5135-11.2012> PMID: 22357858
21. Ahmadian Y, Fumarola F, Miller KD. Properties of networks with partially structured and partially random connectivity. *Physical Review E*. 2015; 91:012820. <https://doi.org/10.1103/PhysRevE.91.012820> PMID: 25679669
22. Mastrogiuseppe F, Ostojic S. Linking connectivity, dynamics, and computations in low-rank recurrent neural networks. *Neuron*. 2018; 99:609–623. <https://doi.org/10.1016/j.neuron.2018.07.003> PMID: 30057201
23. Vu VH. Spectral norm of random matrices. In: *Proceedings of the Thirty-Seventh Annual ACM Symposium on Theory of Computing*. STOC'05. New York, NY, USA: Association for Computing Machinery; 2005. p. 423–430. Available from: <https://doi.org/10.1145/1060590.1060654>.
24. Roy O, Vetterli M. The effective rank: A measure of effective dimensionality. In: *2007 15th European Signal Processing Conference*; 2007. p. 606–610. Available from: <https://ieeexplore.ieee.org/abstract/document/7098875>.
25. Lübke J, Markram H, Frotscher M, Sakmann B. Frequency and dendritic distribution of autapses established by layer 5 pyramidal neurons in the developing rat neocortex: comparison with synaptic innervation of adjacent neurons of the same class. *Journal of Neuroscience*. 1996; 16:3209–3218. <https://doi.org/10.1523/JNEUROSCI.16-10-03209.1996> PMID: 8627359
26. Yin L, Zheng R, Ke W, He Q, Zhang Y, Li J, et al. Autapses enhance bursting and coincidence detection in neocortical pyramidal cells. *Nature Communications*. 2018; 9:4890. <https://doi.org/10.1038/s41467-018-07317-4> PMID: 30459347
27. Morrison A, Aertsen A, Diesmann M. Spike-timing-dependent plasticity in balanced random networks. *Neural Computation*. 2007; 19:1437–1467. <https://doi.org/10.1162/neco.2007.19.6.1437> PMID: 17444756
28. Soltoggio A, Steil JJ. Solving the distal reward problem with rare correlations. *Neural Computation*. 2013; 25:940–978. https://doi.org/10.1162/NECO_a_00419 PMID: 23339615
29. Tao T. Outliers in the spectrum of iid matrices with bounded rank perturbations. *Probability Theory and Related Fields*. 2013; 155:231–263. <https://doi.org/10.1007/s00440-011-0397-9>
30. Aljadeff J, Stern M, Sharpee T. Transition to chaos in random networks with cell-type-specific connectivity. *Physical Review Letters*. 2015; 114:088101. <https://doi.org/10.1103/PhysRevLett.114.088101> PMID: 25768781
31. Zhao S, Sheibanie AF, Oh M, Rabbah P, Nadim F. Peptide neuromodulation of synaptic dynamics in an oscillatory network. *Journal of Neuroscience*. 2011; 31:13991–14004. <https://doi.org/10.1523/JNEUROSCI.3624-11.2011> PMID: 21957260
32. Kruglikov I, Rudy B. Perisomatic GABA release and thalamocortical integration onto neocortical excitatory cells are regulated by neuromodulators. *Neuron*. 2008; 58(6):911–924. <https://doi.org/10.1016/j.neuron.2008.04.024> PMID: 18579081
33. Baimoukhametova DV, Hewitt SA, Sank CA, Bains JS. Dopamine modulates use-dependent plasticity of inhibitory synapses. *Journal of Neuroscience*. 2004; 24:5162–5171. <https://doi.org/10.1523/JNEUROSCI.4979-03.2004> PMID: 15175385
34. Manzke T, Dutschmann M, Schlaf G, Mörschel M, Koch UR, Ponimaskin E, et al. Serotonin targets inhibitory synapses to induce modulation of network functions. *Philosophical Transactions of the Royal Society B: Biological Sciences*. 2009; 364:2589–2602. <https://doi.org/10.1098/rstb.2009.0068> PMID: 19651659
35. Pfeffer CK, Xue M, He M, Huang ZJ, Scanziani M. Inhibition of inhibition in visual cortex: the logic of connections between molecularly distinct interneurons. *Nature Neuroscience*. 2013; 16:1068–1076. <https://doi.org/10.1038/nn.3446> PMID: 23817549
36. Jiang X, Shen S, Cadwell C, Berens P, Sinz F, Ecker A, et al. Principles of connectivity among morphologically defined cell types in adult neocortex. *Science*. 2015; 350(6264):aac9462. <https://doi.org/10.1126/science.aac9462> PMID: 26612957
37. Tsodyks MV, Skaggs WE, Sejnowski TJ, McNaughton BL. Paradoxical effects of external modulation of inhibitory interneurons. *Journal of Neuroscience*. 1997; 17:4382–4388. <https://doi.org/10.1523/JNEUROSCI.17-11-04382.1997> PMID: 9151754
38. Gao P, Ganguli S. On simplicity and complexity in the brave new world of large-scale neuroscience. *Current Opinion in Neurobiology*. 2015; 32:148–155. <https://doi.org/10.1016/j.conb.2015.04.003> PMID: 25932978

39. Finkelstein A, Fontolan L, Economo MN, Li N, Romani S, Svoboda K. Attractor dynamics gate cortical information flow during decision-making. *Nature Neuroscience*. 2021; 24:843–850. <https://doi.org/10.1038/s41593-021-00869-7> PMID: 33875892
40. Motta A, Berning M, Boergens KM, Staffler B, Beining M, Loomba S, et al. Dense connectomic reconstruction in layer 4 of the somatosensory cortex. *Science*. 2019; 366 (6469). <https://doi.org/10.1126/science.aay3134> PMID: 31649140
41. Litwin-Kumar A, Doiron B. Formation and maintenance of neuronal assemblies through synaptic plasticity. *Nature Communications*. 2014; 5:5319. <https://doi.org/10.1038/ncomms6319> PMID: 25395015
42. Zenke F, Agnes EJ, Gerstner W. Diverse synaptic plasticity mechanisms orchestrated to form and retrieve memories in spiking neural networks. *Nature Communications*. 2015; 6:6922. <https://doi.org/10.1038/ncomms7922> PMID: 25897632
43. Bono J, Clopath C. Modeling somatic and dendritic spike mediated plasticity at the single neuron and network level. *Nature Communications*. 2017; 8:706. <https://doi.org/10.1038/s41467-017-00740-z> PMID: 28951585
44. Agnes EJ, Luppi AI, Vogels TP. Complementary inhibitory weight profiles emerge from plasticity and allow flexible switching of receptive fields. *Journal of Neuroscience*. 2020; 40:9634–9649. <https://doi.org/10.1523/JNEUROSCI.0276-20.2020> PMID: 33168622
45. Gütig R, Sompolinsky H. The tempotron: a neuron that learns spike timing–based decisions. *Nature Neuroscience*. 2006; 9:420–428. <https://doi.org/10.1038/nn1643> PMID: 16474393
46. Hendin O, Horn D, Tsodyks MV. The role of inhibition in an associative memory model of the olfactory bulb. *Journal of Computational Neuroscience*. 1997; 4:173–182. <https://doi.org/10.1023/A:1008895429790> PMID: 9154523
47. Song S, Miller KD, Abbott LF. Competitive Hebbian learning through spike-timing-dependent synaptic plasticity. *Nature Neuroscience*. 2000; 3:919–926. <https://doi.org/10.1038/78829> PMID: 10966623