The computational and learning benefits of Daleian neural networks

Anonymous Author(s)
Affiliation
Address
email

Abstract

Dale's principle implies that biological neural networks are composed of neurons that are either excitatory or inhibitory. While the number of possible architectures of such Daleian networks is exponentially smaller than the number of non-Daleian ones, the computational and functional implications of using Daleian networks by the brain are mostly unknown. Here, we use models of recurrent spiking neural networks and rate-based ones to show, surprisingly, that despite the structural limitations on Daleian networks, they can approximate the computation performed by non-Daleian networks to a very high degree of accuracy. Moreover, we find that Daleian networks are more functionally robust to synaptic noise. We then show that unlike non-Daleian networks, Daleian ones can learn efficiently by tuning of single neuron features, nearly as well as learning by tuning individual synaptic weights. Importantly, this suggests a simpler and more biologically plausible learning mechanisms. We therefore suggest that in addition to architectural simplicity, Dale's principle confers computational and learning benefits for biological networks, and offer new directions for constructing and training biologically-inspired artificial neural networks.

1 Introduction

The synaptic connectivity of neurons in the brain is shaped by a variety of physical, chemical, and biological factors [1, 2], even before learning comes into play. For example, the location of neurons and their morphology induce distance-dependent connectivity [3], whereas molecular lock-and-key mechanisms induce cell-type specific connectivity patterns [4]. The functional and computational implications of such constraints - from the architecture of connectomes to cell types - are fundamental to our understanding how real neural circuits are built and function.

A particularly stringent constraint on the design of biological neural networks is imposed by Dale's principle, which asserts that the axonal branches of a neuron release the same set of neurotransmitter molecules [5, 6]. While some exceptions to Dale's principle have been reported [7–10], the implied identity of neurons as either excitatory or inhibitory has been central to our understanding of neural circuit structure and function [11–14]. Moreover, the balance between excitatory and inhibitory populations has been at the core of many computational models of the design of neural circuits and their function [15].

It is clear that Dale's principle has far-reaching implications for the architecture of neural circuits and thereby their computational capacity and learning: For directed networks of N neurons, there are $N \cdot (N-1)$ possible synapses, and so for the general case in which each synapses can be either excitatory, inhibitory, or non-existent (independently of one another) - there are $3^{N\cdot(N-1)}$ possible signed networks connecting the neurons. However, only $\left(2 \cdot 2^{N-1} - 1\right)^N$ of these networks abide by Dale's principle (see SI). Thus, already for 15 neurons, non-Daleian networks outnumber the Daleian

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

ones by a factor of $\sim 10^{32}$. The richness of synaptic connectivity maps that non-Daleian networks offer, the staggering difference in their number compared to Daleian ones, and their potentially higher plasticity (the ability to change synaptic signs in learning) implies that non-Daleian networks may be able to carry computations that Daleian networks cannot. However, Daleian networks are clearly simpler in terms of biological design and development, since the number of neurotransmitters and of receptors expressed by each neuron are more limited, compared to the expected requirements of regulation of a connections in non-Daleian ones. Moreover, since during learning synapses cannot change their sign, learning in Daleian networks is limited to a significantly smaller search space, which could make it easier for Daleian architectures to find networks that perform a desired function. The obvious question is then: what is the brain "missing" by imposing Dale's principle on biological neural networks – if at all?

Analysis of the statistical properties of Daleian connectivity matrices in terms of their eigenvalue spectra suggested that the variance of synaptic weights in this case shapes network dynamics [16], and described the covariance structure of neural activity in such networks [17, 18]. One functional analysis of networks violating Dale's principle [19] suggested they may allow for robust balanced network states using fewer neurons than Daleian networks. Studies of training neural networks that obey Dale's principle suggested that feed-forward artificial neural networks (ANNs) with separate inhibitory and excitatory units can reach similar performance as feed-forward multi-layered perceptrons that are not sign specific [20]. For spiking neural network models, a scheme for constructing networks that are compatible with biophysical constraints on connectivity and perform tasks that are solved by biological neural circuits, has been suggested in [21].

Here, we study the implications of Dale's principle on the function of recurrent neural networks, and their ability to learn. We simulate large ensembles of Daleian and non-Daleian networks and use the response distributions of individual networks, to rich sets of stimuli, to measure the functional similarity between them in terms of the overlap of their spiking response. Surprisingly, we find that almost all non-Daleian networks have a Daleian "neighbor" that is close in functional space, namely it responds in a similar way to the same stimulus. Thus, the computations implemented by non-Daleian networks can be accurately approximated by Daleian ones. We analyze the sensitivity of networks to synaptic perturbations, and find that computations performed by Daleian networks are more robust to random synaptic noise. We further show Daleian networks can learn to approximate the computation that arbitrary non-Daleian networks perform using a simpler and more biologically plausible learning mechanism. Overall, our results suggest that Daleian networks are beneficial both from both a design perspective and a computational one, and offer potential venues for exploring biologically-inspired design of artificial neural networks.

2 Results

To characterize the functional differences between Daleian and non-Daleian networks, we study the encoding of stimuli by these two kinds of networks, using two classes of models of recurrent neural networks - networks of spiking neurons and networks of rate-based neurons. We compare the responses of Daleian networks and non-Daleian ones for the same stimulus or class of stimuli, and explore the nature and implications of their differences.

2.1 Daleian and non-Daleian networks tile the space of neural response distributions in a similar manner

We start by considering the case of networks of N spiking neurons, whose synaptic connectivity is given by an $N \times N$ real matrix, W, where W_{ij} is the synapse from neuron i to neuron j, and $W_{ii} = 0$ for all i (i.e., no self-synapses). The stimulus to the network is given by a vector $\mathbf{s} \in R^N$, such that s_i is a time-independent stimulus to neuron i. The biases of individual neurons towards spiking are represented as $\boldsymbol{\theta} \in R^N$, such that θ_i sets the baseline activity level of neuron i. The activity of the network in a time bin t is represented as a binary vector $\mathbf{x} \in \{0,1\}^N$, such that $x_i = 1$ if neuron i spiked in that time bin, and $x_i = 0$. The dynamics of the network is given by

$$P_{W,\boldsymbol{\theta}}\left(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{s}\right) = \frac{1}{Z} \exp\left(\mathbf{x}_{t-1}^T \cdot W \cdot \mathbf{x}_t + \boldsymbol{\theta} \cdot \mathbf{x}_t + \mathbf{s} \cdot \mathbf{x}_t\right) \tag{1}$$

so the probability of neuron j to spike in time t is given by a sigmoidal function of the sum of its synaptic inputs $\sum_i W_{ij} \cdot x^i_{t-1}$, its bias towards spiking, θ_j, and the stimulus it receives, s_j, where Z is the partition function. Equation (1) defines a transition matrix between the 2^N binary states of the network, M, whose entries are all larger than zero, and so this Markov chain converges to a unique stationary distribution over the states of the network for a given stimulus $\pi_{W,\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s})$. This unique stationary distribution satisfies $\pi_{W,\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s}) \cdot M = \pi_{W,\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s})$, or $\pi_{W,\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s}) \cdot (M - I) = 0$, where I is the identity matrix. Since $\pi_{W,\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s})$ is a vector from the left nullspace of $(M - I)$, and we know that this nullspace has dimension 1, $\pi_{W,\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s})$ can be uniquely found by singular value decomposition (SVD) of $(M - I)$ (see SI).

To compare Daleian and non-Daleian networks, we computed the stationary distributions $\pi_{W,\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s})$ of 5000 Daleian networks and 5000 non-Daleian networks of $N = 10$ neurons, each presented with 100 randomly selected stimuli (Fig. 1a-c). For the non-Daleian (nD) networks, synapses were sampled from a Gaussian distribution $W^{\text{nD}}_{ij} \sim \mathcal{N}(0, \frac{1}{\sqrt{N}})$, where $i, j \in 1 \ldots N$. For the Daleian (D) networks, half of the neurons were selected to be excitatory and half inhibitory, and the outgoing synaptic weights of each neuron were sampled from a half-normal distribution, $W^{\text{D}}_{ij} \sim |\mathcal{N}(0, \frac{1}{\sqrt{N}})| \cdot sign(i)$. The stimuli were sampled from a normal distribution, $s_i \sim \mathcal{N}(0, 1)$ $(i \in 1 \ldots N)$, as in this range the network responses were not dominated by the stimulus or by the recurrent activity in the network, but by the combination of both (see SI). Importantly, while the stationary distributions are an asymptotic property of the model, the rate of convergence of these Markov chains for our networks is very fast, approaching $\pi_{W,\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s})$ at an exponential rate (the rate is governed by the spectral gap of M, which here corresponds to about 5 time steps, or $100\ ms$; see SI). We note that the dynamics defined by eq. (1) converge to $\pi_{W,\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s})$ *irrespective* of the initial conditions of the network, which justifies the comparison of the networks by the stationary distributions of their responses to the same stimulus.

We quantified the functional similarity between all pairs of networks within and between the two ensembles, using the dissimilarity of the stationary response distributions:

$$D_{func}\left(W_k, W_l | \mathbf{s}\right) = D_{JS}\left[\pi_{W_k,\boldsymbol{\theta}_k}(\mathbf{x}|\mathbf{s}) || \pi_{W_l,\boldsymbol{\theta}_l}(\mathbf{x}|\mathbf{s})\right]. \tag{2}$$

where D_{JS} is the Jensen-Shannon divergence – a bounded and symmetric measure of the overlap of the distributions [23]. Figure 1d shows a sub-sample of the resulting $10,000 \times 10,000$ dissimilarity matrix between all networks in the Daleian and non-Daleian ensembles for one such stimulus. The dissimilarity matrices for different stimuli were highly correlated (see SI). The structure of the dissimilarity matrix suggests that Daleian networks do not inhabit a specific part of functional space, but rather that their response distributions tile the space of stationary distributions in a similar way to that of the non-Daleian networks. This is further demonstrated by the 2-dimensional embedding of the dissimilarity matrix (Fig. 1e), which does not show any obvious separation between Daleian and non-Daleian networks. Furthermore, the profiles of distances from a randomly chosen Daleian or non-Daleian network to the other networks are consistent over different choices of networks, as shown in Fig. 1e,f - reflecting that for a typical non-Daleian network, there is a Daleian network which is functionally close to it. We further computed for each stimulus the percent of non-Daleian networks whose nearest network over the two ensembles was a Daleian network, and found that, across all stimuli, the probability that the nearest network to a non-Daleian network would be Daleian is approximately 50%, reflecting, again, the Daleian networks cover the space of response distributions in a similar manner to that of non-Daleian ones. Importantly, repeating the analysis for the relations between networks, averaged over the set of stimuli, namely the matrix given by $\langle D_{func}\left(W_k, W_l | \mathbf{s}\right)\rangle_{\mathbf{s}}$ (where $\langle\rangle_{\mathbf{s}}$ denotes averaging over all stimuli) gave similar structure in terms of the "coverage" of the same space (see SI).

Thus, the results for our sampled networks and stimuli suggest that for a given non-Daleian network, there is a Daleian network that is nearby in terms of its function. To verify that this is not the result of the details of the spiking neural network models that we used, our sampling of the space of non-Daleian and Daleian networks, or the size of the networks - we expand the classes of models we study and their size, and try to learn Daleian approximations to non-Daleian networks.

2.2 Daleian networks learn to approximate non-Daleian ones with high accuracy

We asked how accurately we can approximate the response of a randomly selected non-Daleian network with synaptic connectivity W^{nD} to stimulus $\mathbf{s}$, using a Daleian network W^D, by optimizing its synaptic weights. We started from a randomly chosen Daleian connectivity matrix W^D_{init} of contin-

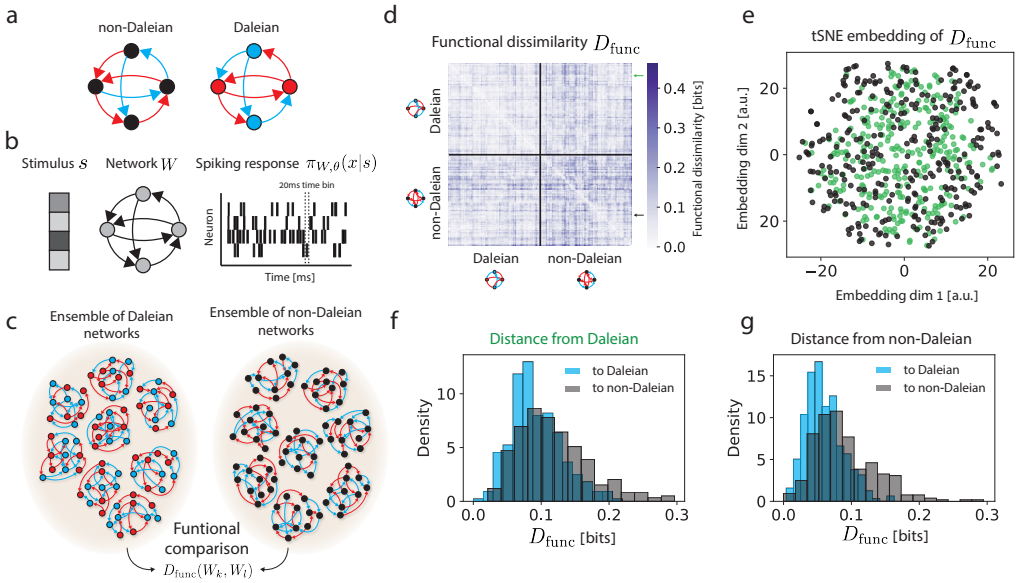

Figure 1: **Daleian and non-Daleian recurrent neural networks cover the same functional space.** (**a**) In non-Daleian networks, the same neuron can have both excitatory and inhibitory outgoing synapses. In Daleian networks, all the synapses going out of a neuron are either exclusively excitatory (red neurons) or inhibitory (blue neurons). (**b**) We simulate the responses of networks of spiking neurons whose connectivity is given by W to a large set of random stimuli, $\{\mathbf{s}_k\}$, and compute the stationary response distribution to each stimulus $\pi_{W,\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s}_k)$. (**c**) We sample 10,000 Daleian and non-Daleian networks, and compute the response distribution of all networks to the same stimuli. We then measure the functional dissimilarity between networks by the overlap of these distributions. (**d**) An example of the pairwise functional dissimilarity matrix between 300 Daleian and 300 non-Daleian networks presented with the same stimulus (only 600 networks are shown for visualization purpose). Rows and columns within the block of Daleian vs. Daleian networks (and non-Daleian vs. non-Daleian ones), were each ordered by hierarchical clustering, see SI [22]. The black horizontal and vertical black lines mark the transition from Dale to non-Dale networks. (**e**) A 2-dimensional tSNE embedding of the dissimilarity matrix from (d). Each dot marks one network, with Daleian networks shown in green, and non-Daleian shown in black. (**f**) A representative example of the distributions of functional dissimilarity values between one randomly chosen Daleian (marked by a green arrow in (d)) network and the rest of the networks in the ensemble. (**g**) Similar to (f) but for a randomly selected non-Daleian network (black arrow box in (d)). Distance profiles shown in (f) and (g) are typical for all networks in the ensemble (SI).

uous synaptic weights, where half of the cells were randomly chosen to have only positive outgoing synapses and half had only negative outgoing synapses, and the diagonal terms are all zero (Fig. 2a). We then used gradient descent on W_D to minimize $\mathcal{L}(W^D) = D_{JS}\left[\pi_{W^D,\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s})||\pi_{W^{nD},\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s})\right]$ (keeping the signs of all the outgoing synapses of each neuron).

We performed this learning for 2000 different non-Daleian networks W^{nD}, each with a different stimulus $\mathbf{s}$, yielding 2000 Daleian approximations (see SI for optimization details). We find that for a random non-Daleian network A there exists, with high probability, a Daleian network B that is orders-of-magnitude closer than a typical non-Daleian network C. Over the whole set, the optimized Daleian networks were two orders-of-magnitude closer to the original network compared to the typical functional distance between networks in our sampled ensembles above, and one order-of-magnitude closer than the nearest network in the ensemble (Fig. 2b). On average, the dissimilarity of an optimized Daleian network to a random non-Daleian responding to an arbitrary stimulus reached a D_{JS} value of $\sim 5 \cdot 10^{-4}\, bits$, which implies that 2000 time bins would be needed to distinguish between the responses of the target non-Daleian network and the optimized Daleian one. (For reference, optimizing a non-Daleian network to approximate a non-Daleian network reached D_{JS} values of $\sim 3 \cdot 10^{-5}\, bits$).

We next asked how well can an optimized Daleian network approximate a non-Daleian one on a set of stimuli, i.e. how similar they are in terms of the function they compute. We therefore repeated the learning described above, but this time the goal was to minimize $\langle \mathcal{L}(W^D) \rangle_{\mathbf{s}}$. Again, the optimized Daleian networks were highly accurate in approximating the non-Daleian ones, now reaching an average D_{JS} value of $\sim 6 \cdot 10^{-3}\, bits$ when training only W^D or when training both W^D and θ (Fig. 2c). These differences mean it would take about 200 time bins to tell apart the responses of the Daleian network from the non-Daleian one. Again, if we consider time bins on the order of $20\, ms$, this means that a non-Daleian network and its Daleian approximation would be distinguishable on the set of stimuli on a timescale of more than $\sim 4000\, ms$ of neural activity.

We emphasize that we did not optimize the cell types of a Daleian network that learns to approximate a non-Daleian one (i.e. we did not pick which neurons would be excitatory and which inhibitory in each network) - which means that our Daleian approximations to the non-Daleian ones are an upper bound on the ability of Daleian networks to approximate non-Daleian ones, and also reflect how capable Daleian. Interestingly, the learned networks had significantly broader distributions of synaptic weights, and Daleian optimized networks were significantly sparser than non-Daleian optimized networks (see SI).

To give a more biologically relevant interpretation to these differences, we asked what level of synaptic noise would result in similar functional dissimilarity values. Thus, for randomly sampled network W^{nD} and stimulus $\mathbf{s}$, we computed the Jensen-Shannon divergence between the stationary distribution of the original W^{nD} and noisy versions of it $W^{nD}_{\epsilon} = W^{nD} \odot \epsilon$, for different levels of multiplicative synaptic noise ϵ (see SI). Averaging over different noise realizations, networks, and stimuli - we found that a functional dissimilarity value of $D_{JS}[\pi_{W^{nD},\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s})||\pi_{W^{nD}_{\epsilon},\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s})] \sim 5{\cdot}10^{-3}$ bits, a typical value for the optimized Daleian approximation, corresponds to a synaptic noise level of $\sim 25\%$ (Fig. 2d). Given the variability of activation strength of the same synaptic connections in real neural networks, and their dynamic nature, these results suggest that Dale networks can approximate non-Daleian ones very close to the possible limit set by synaptic and neuronal noise [24].

Next, we studied how well Daleian networks can approximate non-Daleian ones using another commonly studied class of network models – continuous-time recurrent firing-rate networks [25] – which also allowed us to extend our analysis to much larger networks. Here. neural activity is described by firing rates of the N neurons, which are denoted by $\mathbf{r} \in R^N$. Correspondingly, the dynamics is given by a differential equation with integration time constant τ, and so the response of the network to a stimulus $\mathbf{s}$, is given by

$$\tau \frac{d\mathbf{r}}{dt} = -\mathbf{r} + W^T \mathbf{r} + \mathbf{s}\,. \tag{3}$$

and the steady state solution is given by $\mathbf{r}_{ss} = \left(I - W^T\right)^{-1} \mathbf{s}$. We then asked, again, whether we can approximate the computation performed by an arbitrary network W^{nD} using a Daleian network W^D, but this time for networks of $N = 100$ neurons. We repeated the analysis we performed for the spiking networks on the rate-based one, but here for each randomly chosen non-Daleian network we used gradient descent to minimize a loss function given by $\mathcal{L}(W^D) = \left\| \left(I - \left(W^{nD}\right)^T\right)^{-1} - \left(I - \left(W^D\right)^T\right)^{-1} \right\|^2$. We found that while the optimized Daleian matrix W^D is very different from the W^{nD} of the non-Daleian network it approximates, it gives a very similar stimulus-to-steady-state mapping: We computed the mean Pearson correlation between the steady states of W^{nD} and those predicted by the Daleian approximation W^D for 1000 different stimuli, and found that the typical Pearson correlation between them is $\sim 99\%$. Critically, the resulting correlation values were very similar to the accuracy of learning to approximate a non-Daleian networks by other non-Daleian networks (Fig. 2e).

We conclude that non-Daleian networks can be approximated with high accuracy by Daleian ones. Moreover, since our learning of the optimized Daleian approximations is based on simple gradient descent and randomly chosen initial Daleian network structure, it is clear that our results are (again) are an upper-bound on how close one can get to a non-Daleian network with a Daleian one. Having established that there is no clear functional loss in using Daleian networks compared to the supposedly more powerful non-Daleian ones, we turn to explore the sensitivity and learning dynamics of Daleian and non-Daleian networks, and seek functional benefits in using Daleian networks.

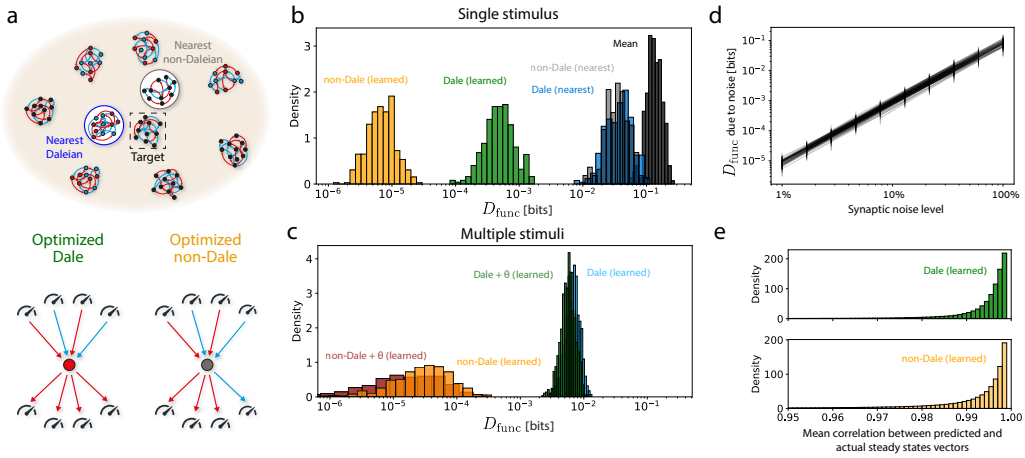

Figure 2: **Daleian networks can learn to accurately approximate the function of non-Daleian networks.** (**a**) For a given non-Daleian network W^{nD} and stimulus **s**, we compute the distance to the nearest Daleian network in the sampled ensemble (blue), the nearest non-Daleian network in the sampled ensemble (gray). We also optimize a random Daleian network to approximate the target non-Daleian network (these learned networks will be later denoted in green) and learn a non-Daleian network that is optimized to approximate the non-Daleian network (later shown in orange). In the last two cases, learning is done by gradient descent on all synapses, see text. (**b**) The histograms of distances of the sampled and learned networks from 2000 different non-Daleian networks and 2000 different stimuli. Different colors show the distances to the target non-Daleian networks from the different classes of networks; color as described in (a). For reference, we also show a histogram of the mean distance from each non-Daleian network to all other networks in the ensemble for the corresponding stimulus (black). (**c**) Functional dissimilarity between the stationary distribution of a network W^{nD} responding to stimulus **s**, and noisy variations of W^{nD} for different levels of synaptic noise (ranging from 1% to 100%; noise is independent for each synapse). Error bars correspond to standard deviations over 100 different noise realizations. Results shown for 30 different networks and 30 different stimuli. (**d**) Empirical cumulative distribution function of synaptic weights before learning (black), and after learning for both Daleian networks (green) and non-Daleian networks (orange). (**e**) Accuracy of the Daleian approximation (green) and a non-Daleian approximation (orange) for the continuous-time recurrent firing-rate model and $N = 100$ neurons, measured by the correlation of their steady state activity vectors.

2.3 Robustness and adaptability of Daleian networks are superior to non-Daleian ones

In learning and optimizing their function, neural networks need to balance changes that improve their performance while retaining their previously learned function. This suggests that certain network changes would result in altered function, whereas others should not. To explore the "design" of learnability of non-Daleian and Daleian networks and their robustness, we next asked how sensitive they are to changes in their synaptic connectivity maps and to the properties of individual neurons.

We first explored the effect of changes in synaptic weights on the function of a given network by computing $f_W(\Delta W) = D_{KL}[\pi_{W,\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s})||\pi_{W+\Delta W,\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s})]$ and then the Hessian of f, given by the matrix $[H_f]^W_{ij,kl} = \frac{\partial^2 f}{\partial W_{ij}\partial W_{kl}}$. We note that in this case, the Hessian of H_f (evaluated at W) is actually the Fisher Information Matrix of $\pi_{W,\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s})$, which provides a principled interpretation of the analysis and the results that follow [26]. The Hessian measures the curvature of f in all directions of parameter space, where each direction corresponds to changing the weight of a single synapse. The trace of the H then measures the overall "flatness" of f around W. From a biological perspective, the flatness corresponds to the robustness of the network to synaptic fluctuations. We thus computed $\mathrm{Tr}\,[H_f] = \sum_{i,j\in 1\ldots N, i\neq j} \frac{\partial^2 f}{\partial W_{ij}^2}$, for each of the networks in the Daleian and non-Daleian ensembles, and for each sampled stimulus. To compare the sensitivities of Daleian and non-Daleian networks, we computed the ratio between the average trace of Daleian networks and that of non-Daleian networks, separately for each stimulus. This ratio quantifies whether Daleian network were more sensitive (ratio

> 1) or less sensitive (ratio < 1) to random synaptic changes compared to non-Daleian networks. We found that for the 98% of the stimuli, Daleian networks were *less* sensitive than non-Daleian ones to perturbations of synaptic weights (Fig. 3a). Given the inherent noise in biological synapses, such robustness carries a clear functional advantage.

Increased robustness to synaptic fluctuations may hinder the ability of a network to adjust its reactions to changes in the environment. We therefore repeated the sensitivity analysis above, but this time computed the Hessian with respect to changes in the bias of individual neurons θ_i, $[H_f]^{\boldsymbol{\theta}}_{ij} = \frac{\partial^2 f}{\partial\theta_i \partial\theta_j}$ and computed $\mathrm{Tr}\,[H_f] = \sum_{i\in 1\ldots N} \frac{\partial^2 f}{\partial\theta_i^2}$, for each of the networks and each stimulus. We then computed the ratio between the average sensitivity of Daleian networks to threshold perturbations, and that of non-Daleian networks, separately for each stimulus. We found that for 84% of the stimuli, Daleian networks were, on average, *more* sensitive to changes in the threshold parameters compared to non-Daleian ones (Fig. 3b).

Thus, the differences in sensitivities of the two classes of networks mean that Daleian networks are more resilient to random synaptic fluctuations, and that they may learn effectively by tuning the baseline firing rates of individual neurons.

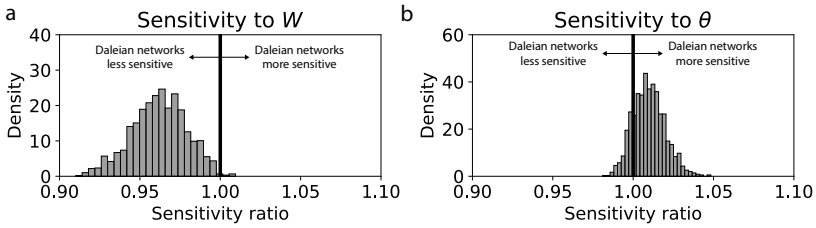

Figure 3: **Daleian networks are less sensitive to synaptic fluctuations and more sensitive to changes of single neuron parameters.** (**a**) Histogram of the ratio of the average sensitivity to synaptic perturbation of Daleian and the average sensitivity of non-Daleian networks, shown for 1000 different stimuli. Average sensitivity is measured using the trace of each of the hessian matrices of 1000 networks for a given stimulus. (**b**) Similar to (a) but sensitivity is computed with respect to neuronal biases vector $\boldsymbol{\theta}$.

2.4 Daleian networks learn efficiently by tuning only single neuron properties

We therefore asked then how learning through changing of individual neuronal properties rather than individual synapses would work for Daleian and non-Daleian networks. Given the experimentally observed homeostatic scaling of synaptic weights [27] and the identification of "architectural" features that shape the function of neural networks [28, 29], it is clear that learning by adapting neuronal features would have clear biological benefits and is more plausible than learning by general changes over the whole connectivity matrix. We thus compared the results of learning by gradient descent over $\{W_{ij}\}$, to that of learning by more macroscopic, neuron-level features. Specifically, we used a scaling factor for all the outgoing synapses of a neuron $\lambda_i^{out} > 0$, a scaling factor for all the incoming synapses of a neuron $\lambda_i^{in} > 0$ and the firing bias θ_i. In this learning mechanism, each synapse is re-scaled according to the pre-synaptic and post-synaptic scaling factors $W_{ij} \to \lambda_i^{out}\lambda_j^{in} \cdot W_{ij}$ (λ_i^{out} and λ_i^{in} are initialized to 1 before learning), and each neuron's bias is adapted individually $\theta_i \to \theta_i + \Delta\theta_i$ ($\Delta\theta_i$ is initialized to 0 before learning). This means that instead of learning by coordinated changes to the $N(N-1)$ terms of the full synaptic connectivity map, we learn by changing $3N$ parameters, which correspond to known biological features at the level of individual neurons (Fig. 4a). Clearly, synaptic learning is more potent than the neuron-level learning – for a given network, any learning trajectory (in weights-space) due to neuronal learning is also achievable by synaptic learning, but not vice versa. However, as we shall see, they act very differently in terms of the accuracy of synaptic and neuron-level learning on Daleian vs. non-Daleian networks.

We further asked how well the different classes of networks would approximate a natural response distribution of biological neural networks. We used real spiking patterns of groups of 10 neurons recorded in macaque monkeys performing a visual classification task [30], to compare the accuracy of

synaptic learning and neuron-level learning in mapping a random stimulus to the empirical response distributions of real neurons. We found that while both Daleian and non-Daleian networks achieved higher accuracy using synaptic learning (Fig. 4b), learning by changing single neuron properties achieved higher performance for Daleian networks compared to non-Daleian networks (Fig. 4c,d). This suggests that the constraints set by Dale's principle change the relative effectiveness of these learning mechanisms.

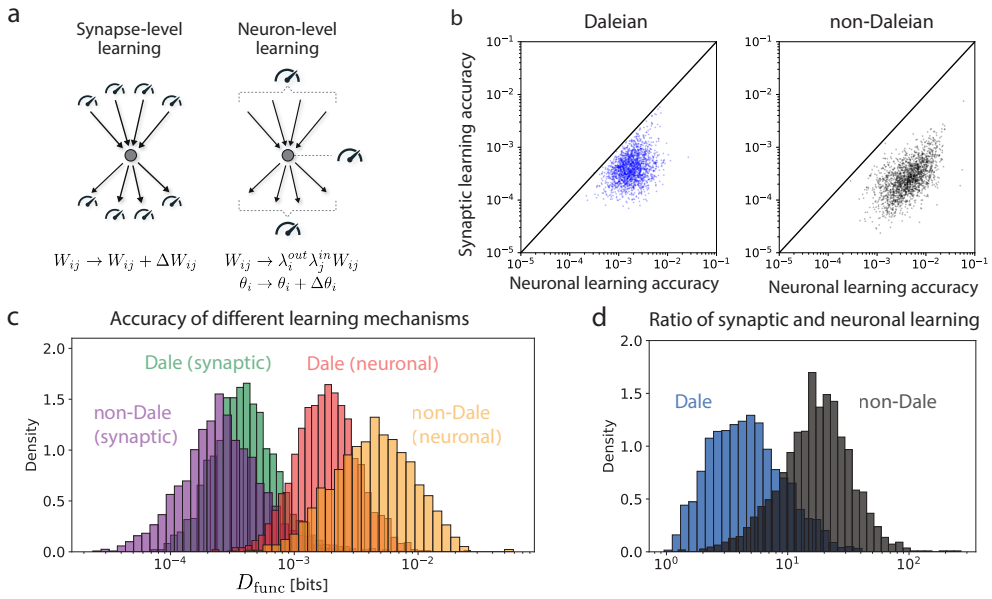

Figure 4: **Daleian networks can learn efficiently by tuning single-neuron properties.** (**a**) In synaptic learning (left), each synapse is separately tuned, while in neuronal, changes to neuronal properties induce correlated across synapses (see text). (**b**) Accuracy of synaptic learning is compared to neuron-level learning, for both Daleian and non-Daleian networks. Each point is the mean D_{JS} when learning the mapping from a random stimulus $\mathbf{s}$ to a random target distribution π_{target}, averaged over 30 Daleian (left) or non-Daleian (right) initial networks. Results shown for 2000 different $(\mathbf{s}, \pi_{\text{target}})$ pairs. (**c**) Comparison of the learning accuracy of Daleian and non-Daleian networks using synaptic and neuron-level learning mechanisms. Histograms show the distributions of average D_{JS} between 30 Daleian/non-Daleian networks trained to map a random stimulus s to a target distribution π_{target}, using either synaptic or neuronal learning. (**d**) The ratio between the accuracy of synaptic learning and neuron-level learning for Daleian and non-Daleian networks, as measured by the ratio of mean D_{JS} values from (b).

2.5 Information coding by Daleian networks

As we have shown that Daleian networks are able to accurately approximate the responses of non-Daleian networks for a given stimulus, we next asked how the two classes differ in their encoding of a set of stimuli, which would reflect on the ability of downstream neurons to decode stimuli from network responses. We computed the information that the responses of a network whose synaptic connectivity is given by W, carries about a set of $k = 500$ different stimuli. We first computed the entropy of the stationary distribution for each stimulus, $H\left[\pi_{W,\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s}_i)\right] = -\sum_{\mathbf{x}} \pi_{W,\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s}_i) \cdot \log_2 \pi_{W,\boldsymbol{\theta}}(\mathbf{x}|\mathbf{s}_i)$. The total entropy of the network's activity is given by $H\left[\pi_{W,\boldsymbol{\theta}}\right] = -\sum_{\mathbf{x}} \pi_{W,\boldsymbol{\theta}}(\mathbf{x}) \cdot \log_2 \pi_{W,\boldsymbol{\theta}}(\mathbf{x})$, where $\pi_{W,\boldsymbol{\theta}}(\mathbf{x}) = \frac{1}{k}\sum_{i=1\ldots k} \pi_{W,\boldsymbol{\theta}}(\mathbf{x}|\,\mathbf{s}_i)$, where we assumed a uniform distribution over the set of stimuli. We then computed the mutual information between the set of stimuli and the stationary responses,

$$I_W(\mathbf{s};\mathbf{x}) = H[\pi_{W,\theta}(\mathbf{x})] - \langle H[\pi_{W,\theta}(\mathbf{x}|\mathbf{s}_i)]\rangle_{\mathbf{s}} \tag{4}$$

where $\langle\rangle_{\mathbf{s}}$ denotes averaging over all stimuli. We computed I_W for 1000 Daleian networks and 1000 non-Daleian networks of $N = 10$ neurons, and found that for randomly sampled networks, the responses of Daleian networks were more informative about the stimulus (Fig. 5a). We then

asked whether Daleian and non-Daleian differ in their ability to learn more informative stimulus representations. Using both synaptic and neuronal learning, we trained the networks to maximize the mutual information in eq. (4). We found that Daleian networks reach similar performance using synaptic learning, but remain more informative (on average) when using neuron-level learning (Fig. 5b).

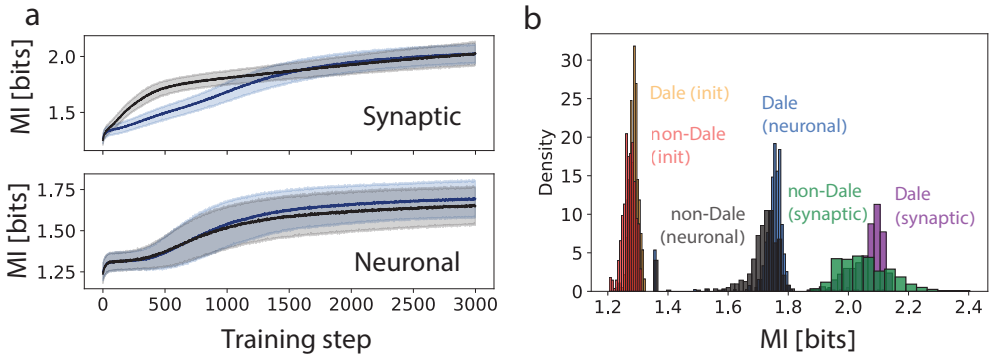

Figure 5: **Random Daleian networks are more informative than non-Daleian ones, whereas trained Daleian networks carry similar information to non-Daleian ones.** (**a**) The mean of the information values carried by 500 Daleian (blue) and 500 non-Daleian (gray) networks, is shown as a function of the training steps. Each network was trained to optimize its information on a given set of 500 stimuli using synaptic learning (top) and neuronal learning (bottom). Error bars represent 1 standard deviation. (**b**) Distribution of information values carried by the networks on a set of 500 random stimuli before learning, with Daleian networks shown in yellow, and non-Daleian networks shown in orange. Distribution of information values of network trained to maximize the carried information by neuronal learning (Daleian in blue, non-Daleian in gray), and using synaptic learning (Dale in Purple, non-Dale in green).

3 Discussion

We found that despite the clear constraint that Dale's principle imposes on the architecture of neural networks, it does not seem to incur a clear cost in terms of the computational capacity of Daleian networks or their ability to learn. We were able to approximate the responses of arbitrary non-Daleian networks to a range of stimuli, with Daleian networks, to a very high degree of accuracy. In fact, this accuracy is so high that it is on par with the level of intrinsic noise level of real biological networks. Moreover, we found that Daleian networks are more robust to synaptic noise, and that they can learn by tuning of neuron-level parameters, nearly as well as can be achieved by general synaptic learning rules. Our results suggest that in addition to the clear structural and developmental benefits of using Daleian neural networks, the brain may actually gain from using them, as they offer similar computational power, higher robustness to noise, and simpler learning mechanisms.

The idea that structural constraints on networks can affect learning trajectories and define the loss surface has a long history in constructing and training artificial neural networks [31–33]. We hypothesize that some of the benefits that Dale's principles confers on recurrent biology-inspired neural networks may extend to artificial ones. It will be especially interesting to explore the parallels of our results in terms of shaping the search space of sign-constrained ANNs, the possibility of efficient learning by neuronal parameters rather than synaptic ones, and the robustness and sensitivity they may offer and the potential for preventing catastrophic forgetting.

Our analysis is limited by our use of dense and relatively small networks, and should be extended to large and sparse networks, and more naturalistic stimuli. Moreover, biological networks have synaptic weights that are distributed log-normally (though this choice did not affect our learned networks).

References

[1] Kosuke Takagi. Energy constraints on brain network formation. *Scientific Reports*, 11(1):1–8, 2021.

[2] Michael W Reimann, Anna-Lena Horlemann, Srikanth Ramaswamy, Eilif B Muller, and Henry Markram. Morphological diversity strongly constrains synaptic connectivity and plasticity. *Cerebral Cortex*, 27(9):4570–4585, 2017.

[3] Bernhard Hellwig. A quantitative analysis of the local connectivity between pyramidal neurons in layers 2/3 of the rat visual cortex. *Biological cybernetics*, 82(2):111–121, 2000.

[4] Joshua R Sanes and S Lawrence Zipursky. Synaptic specificity, recognition molecules, and assembly of neural circuits. *Cell*, 181(3):536–556, 2020.

[5] John Carew Eccles. From electrical to chemical transmission in the central nervous system: the closing address of the sir henry dale centennial symposium cambridge, 19 september 1975. *Notes and records of the Royal Society of London*, 30(2):219–230, 1976.

[6] Piergiorgio Strata, Robin Harvey, et al. Dale's principle. *Brain research bulletin*, 50(5):349–350, 1999.

[7] DW Golding. Patterns of peptide discharge—implications for dale's principle. In *The Peptidergic Neuron*, pages 7–12. Springer, 1996.

[8] Wayne S Sossin, Alejandro Sweet-Cordero, and Richard H Scheller. Dale's hypothesis revisited: different neuropeptides derived from a common prohormone are targeted to different processes. *Proceedings of the National Academy of Sciences*, 87(12):4845–4848, 1990.

[9] D Sulzer and S Rayport. Dale's principle and glutamate corelease from ventral midbrain dopamine neurons. *Amino acids*, 19(1):45–52, 2000.

[10] Erik Svensson, John Apergis-Schoute, Geoffrey Burnstock, Michael P Nusbaum, David Parker, and Helgi B Schiöth. General principles of neuronal co-transmission: insights from multiple model systems. *Frontiers in neural circuits*, page 117, 2019.

[11] Carl Van Vreeswijk and Haim Sompolinsky. Chaos in neuronal networks with balanced excitatory and inhibitory activity. *Science*, 274(5293):1724–1726, 1996.

[12] Tim P Vogels, Kanaka Rajan, and Larry F Abbott. Neural network dynamics. *Annu. Rev. Neurosci.*, 28:357–376, 2005.

[13] Nicolas Brunel. Dynamics of sparsely conntected networks of excitatory and inhibitory spiking neurons. *Journal of Computational Neuroscience*, 8:183–208, 2000.

[14] Michael Okun and Ilan Lampl. Instantaneous correlation of excitation and inhibition during ongoing and sensory-evoked activities. *Nature neuroscience*, 11(5):535–537, 2008.

[15] David C Somers, Sacha B Nelson, and Mriganka Sur. An emergent model of orientation selectivity in cat visual cortical simple cells. *Journal of Neuroscience*, 15(8):5448–5465, 1995.

[16] Kanaka Rajan and Larry F Abbott. Eigenvalue spectra of random matrices for neural networks. *Physical review letters*, 97(18):188104, 2006.

[17] Yu Hu and Haim Sompolinsky. The spectrum of covariance matrices of randomly connected recurrent neuronal networks. *bioRxiv*, 2020.

[18] Jonathan Kadmon and Haim Sompolinsky. Transition to chaos in random neuronal networks. *Physical Review X*, 5(4):041030, 2015.

[19] Victor J Barranca, Asha Bhuiyan, Max Sundgren, and Fangzhou Xing. Functional implications of dale's law in balanced neuronal network dynamics and decision making. *Frontiers in Neuroscience*, 16, 2022.

[20] Jonathan Cornford, Damjan Kalajdzievski, Marco Leite, Amélie Lamarquette, Dimitri Michael Kullmann, and Blake Aaron Richards. Learning to live with dale's principle: Anns with separate excitatory and inhibitory units. In *International Conference on Learning Representations*, 2020.

[21] Brian DePasquale, Mark M Churchland, and LF Abbott. Using firing-rate dynamics to train recurrent networks of spiking model neurons. *arXiv preprint arXiv:1601.07620*, 2016.

[22] Ziv Bar-Joseph, David K Gifford, and Tommi S Jaakkola. Fast optimal leaf ordering for hierarchical clustering. *Bioinformatics*, 17(suppl_1):S22–S29, 2001.

[23] Jianhua Lin. Divergence measures based on the shannon entropy. *IEEE Trans. Information Theory*, 37(1):145–151, 1991.

[24] Noam E Ziv and Naama Brenner. Synaptic tenacity or lack thereof: spontaneous remodeling of synapses. *Trends in neurosciences*, 41(2):89–99, 2018.

[25] Sadra Sadeh and Claudia Clopath. Theory of neuronal perturbome in cortical networks. *Proceedings of the National Academy of Sciences*, 117(43):26966–26976, 2020.

[26] Benjamin B Machta, Ricky Chachra, Mark K Transtrum, and James P Sethna. Parameter space compression underlies emergent theories and predictive models. *Science*, 342(6158):604–607, 2013.

[27] Gina G. Turrigiano and Sacha B. Nelson. Homeostatic plasticity in the developing nervous system. *Nature Reviews Neuroscience*, 5(9):97–107, 2004.

[28] Adam Haber and Elad Schneidman. Learning the architectural features that predict functional similarity of neural networks. *bioRxiv*, 2020.

[29] R Milo, S Shen-Orr, S Itzkovitz, N Kashtan, D Chklovskii, and U Alon. Network motifs: simple building blocks of complex networks. *Science*, 298(5594):824–827, 2002.

[30] Roozbeh Kiani, Christopher J Cueva, John B Reppas, and William T Newsome. Dynamics of neural population responses in prefrontal cortex indicate changes of mind on single trials. *Current Biology*, 24(13):1542–1547, 2014.

[31] Devon Graham, Junhao Wang, and Siamak Ravanbakhsh. Equivariant entity-relationship networks. *arXiv preprint arXiv:1903.09033*, 2019.

[32] Steven J Nowlan and Geoffrey E Hinton. Simplifying neural networks by soft weight-sharing. *Neural computation*, 4(4):473–493, 1992.

[33] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. *Advances in neural information processing systems*, 31, 2018.

Checklist

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes] Limitations of our work are described in the final paragraph of the discussion.
   (c) Did you discuss any potential negative societal impacts of your work? [N/A]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes] We specify the parameters of the models from which the network ensembles were sampled, as well as details about the optimization objective functions and hyperparameters.

(b) Did you include complete proofs of all theoretical results? [Yes] The only theoretical claim in the paper is the ratio between Daleian andn non-Daleian networks of a given size, whose derivation is in the Appendix.

3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] Link to a repository containing the code will be attached as supplemental material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] We quantify and present the variability of our results with respect to different networks, stimuli and noise realizations throughout the work.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] All analysis was conducted on an internal cluster.
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes] We cite the paper describing the neural recordings we have used in Section 2.4.
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [No]
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]