

LEARNING TO LIVE WITH DALE’S PRINCIPLE: ANNS WITH SEPARATE EXCITATORY AND INHIBITORY UNITS

Jonathan Cornford
Mila/McGill

Damjan Kalajdzievski
Mila/UdeM

Marco Leite
UCL

Amélie Lamarquette
University of Cambridge

Dimitri M. Kullmann
UCL

Blake Richards
Mila/McGill

ABSTRACT

The units in artificial neural networks (ANNs) can be thought of as abstractions of biological neurons, and ANNs are increasingly used in neuroscience research. However, there are many important differences between ANN units and real neurons. One of the most notable is the absence of Dale’s principle, which ensures that biological neurons are either exclusively excitatory or inhibitory. Dale’s principle is typically left out of ANNs because its inclusion impairs learning. This is problematic, because one of the great advantages of ANNs for neuroscience research is their ability to learn complicated, realistic tasks. Here, by taking inspiration from feedforward inhibitory interneurons in the brain we show that we can develop ANNs with separate populations of excitatory and inhibitory units that learn just as well as standard ANNs. We call these networks Dale’s ANNs (DANNs). We present two insights that enable DANNs to learn well: (1) DANNs are related to normalization schemes, and can be initialized such that the inhibition centres and standardizes the excitatory activity, (2) updates to inhibitory neuron parameters should be scaled using corrections based on the Fisher Information matrix. These results demonstrate how ANNs that respect Dale’s principle can be built without sacrificing learning performance, which is important for future work using ANNs as models of the brain. The results also may have interesting implications for how inhibitory plasticity in the real brain operates.

1 INTRODUCTION

In recent years, artificial neural networks (ANNs) have been increasingly used in neuroscience research for modelling the brain at the algorithmic and computational level (Richards et al., 2019; Kietzmann et al., 2018; Yamins & DiCarlo, 2016). They have been used for exploring the structure of representations in the brain, the learning algorithms of the brain, and the behavioral patterns of humans and animals (Bartunov et al., 2018; Donhauser & Baillet, 2020; Michaels et al., 2019; Schrimpf et al., 2018; Yamins et al., 2014; Kell et al., 2018). Evidence shows that the ability of ANNs to match real neural data depends critically on two factors. First, there is a consistent correlation between the ability of an ANN to learn well on a task (e.g. image recognition, audio perception, or motor control) and the extent to which its behavior and learned representations match real data (Donhauser & Baillet, 2020; Michaels et al., 2019; Schrimpf et al., 2018; Yamins et al., 2014; Kell et al., 2018). Second, the architecture of an ANN also helps to determine how well it can match real brain data, and generally, the more realistic the architecture the better the match (Schrimpf et al., 2018; Kubilius et al., 2019; Nayebi et al., 2018). Given these two factors, it is important for research that uses ANNs to model the brain, to use ANNs that have as realistic an architecture as possible, but which also learn well (Richards et al., 2019; Kietzmann et al., 2018; Yamins & DiCarlo, 2016).

Although there are numerous disconnects between ANNs and the architecture of biological neural circuits, one of the most notable is the lack of adherence to Dale’s principle, which states that neurons release the same set of neurotransmitters from all of their synaptic terminals (Eccles, 1976). For ANNs, this means that units cannot have a mix of positive and negative output weights, and furthermore, weights cannot change their sign after initialisation. In other words, units can only

be excitatory or inhibitory. Some computational neuroscience studies use ANNs that have separate excitatory and inhibitory units (Song et al., 2016; Ingrosso & Abbott, 2019; Miconi, 2017; Minni et al., 2019), but many do not, particularly those that train on complicated tasks (Bartunov et al., 2018; Donhauser & Baillet, 2020; Michaels et al., 2019; Schrimpf et al., 2018; Yamins et al., 2014; Kell et al., 2018). Why do many neuroscience studies that use ANNs not incorporate Dale’s principle? The reason is that this architectural constraint impairs the ability to learn—a fact that is known to many researchers who have tried to train such ANNs, but one that is rarely discussed in the literature. However, when we seek to compare ANNs to real brains, or use them to explore biologically inspired learning rules (Bartunov et al., 2018; Whittington & Bogacz, 2019; Lillicrap et al., 2020), ideally we would use a biologically plausible architecture with distinct populations of excitatory and inhibitory neurons, and at the same time, *we would still be able to match the learning performance of standard ANNs without such constraints*.

Here, we develop ANN models with separate excitatory and inhibitory units that are able to learn as well as standard ANNs. Specifically, we develop a novel form of ANN, which we call a “Dale’s ANN” (DANN), based on feed-forward inhibition in the brain (Pouille et al., 2009). Our novel approach is different from the standard solution, which is to create ANNs with separate excitatory and inhibitory units by constraining whole columns of the weight matrix to be all positive or negative. Throughout this manuscript, we refer to this standard approach as “ColumnEi” models. We have departed from the ColumnEI approach in our work because it has three undesirable attributes. First, constrained weight matrix columns impair learning because they limit the potential solution space (Amit et al., 1989; Parisien et al., 2008). Second, modelling excitatory and inhibitory units with the same connectivity patterns is biologically misleading, because inhibitory neurons in the brain tend to have very distinct connectivity patterns from excitatory neurons (Tremblay et al., 2016). Third, real inhibition can act in both a subtractive and a divisive manner (Atallah et al., 2012; Wilson et al., 2012; Seybold et al., 2015; Pouille et al., 2013), which may provide important functionality.

Given these considerations, in DANNs, we utilize a separate pool of inhibitory neurons with a distinct, more biologically realistic connectivity pattern, and a mixture of subtractive and divisive inhibition (Fig. 1). This loosely mimics the fast feedforward subtractive and divisive inhibition provided by fast-spiking interneurons in the cortical regions of the brain (Atallah et al., 2012; Hu et al., 2014; Lourenço et al., 2020). In order to get DANNs to learn as well as standard ANNs we also employ two key insights:

- It is possible to view this architecture as being akin to normalisation schemes applied to the excitatory input of a layer (Ba et al., 2016; Ioffe & Szegedy, 2015; Wu & He, 2018), and we use this perspective to motivate DANN parameter initialisation.
- It is important to scale the inhibitory parameter updates based on the Fisher information matrix, in order to balance the impact of excitatory and inhibitory parameter updates, similar in spirit to natural gradient approaches (Martens, 2014).

Altogether, our principle contribution is a novel architecture that obey’s Dale’s principle, and that we show can learn as well as standard ANNs on machine learning benchmark tasks. This provides the research community with a new modelling tool that will allow for more direct comparisons with real neural data than traditional ANNs allow, but which does not suffer from learning impairments. Moreover, our results have interesting implications for inhibitory plasticity, and provide a means for future research into how excitatory and inhibitory neurons in the brain interact at the algorithmic level.

2 BIOLOGICALLY INSPIRED NETWORKS THAT OBEY DALE’S PRINCIPLE

2.1 MODEL DEFINITION

Our design for DANNs takes inspiration from the physiology of feed-forward inhibitory microcircuits in the neocortex and hippocampus. Based on these circuits, and Dale’s principle, we construct DANNs with the following architectural constraints:

1. Each layer of the network contains two distinct populations of units, an excitatory and an inhibitory population.

2. There are far fewer inhibitory units than excitatory units in each layer, just as there are far more excitatory neurons than inhibitory neurons (~ 5 -10 times) in cortical regions of the brain (Tremblay et al., 2016; Hu et al., 2014).
3. As in real neural circuits, only the excitatory populations project between layers, and both the excitatory and inhibitory populations of layer receive excitatory projections from the layer below.
4. All of the synaptic weights are strictly non-negative, and inhibition is enforced via the activation rules for the units (eq. 1).
5. The inhibitory population inhibits the excitatory population through a mixture of subtractive and divisive inhibition.

This constrained architecture is illustrated in Figure 1.

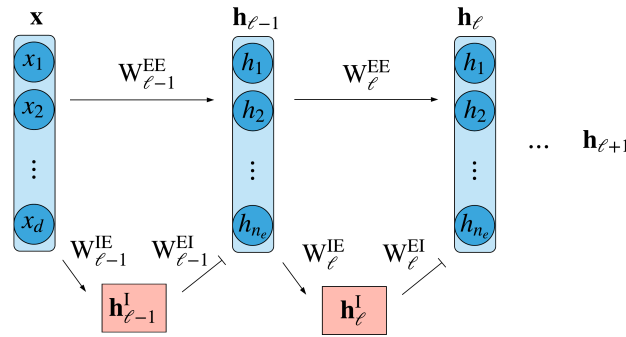


Figure 1: Illustration of DANN architecture. Lines with arrow ends indicate excitatory projections. Lines with bar ends indicate inhibitory projections, which can be *both* subtractive and divisive.

Formally, we define the network as follows. Input to the network is received as a vector of positive scalar values $\mathbf{x} \in \mathbb{R}_+^d$, which we consider to be the first excitatory population. Each hidden layer, ℓ , is comprised of a vector of excitatory units $\mathbf{h}_\ell \in \mathbb{R}_+^{n_e}$ and inhibitory units $\mathbf{h}_\ell^I \in \mathbb{R}_+^{n_i}$, in-line with constraint (1) above. Note, for the first layer ($\ell = 1$), we have $\mathbf{h}_\ell = \mathbf{x}$ and $n_e = d$. Next, based on constraint (2) we set $n_e \gg n_i$, and use 10% inhibitory units as default. Following constraint (3), both the excitatory and inhibitory units receive inputs from the excitatory units in the layer below ($\mathbf{h}_{\ell-1}$), but the inhibitory units do not project between layers. Instead, excitatory units receive inputs from the inhibitory units of the same layer. In-line with constraint (4), we have three sets of strictly non-negative synaptic weights, one for the excitatory connections between layers, $\mathbf{W}_\ell^{EE} \in \mathbb{R}_+^{n_e \times n_e}$, one for the excitatory projection to the inhibitory units $\mathbf{W}_\ell^{IE} \in \mathbb{R}_+^{n_i \times n_e}$, and one for the inhibitory projections within layer $\mathbf{W}_\ell^{EI} \in \mathbb{R}_+^{n_e \times n_i}$. Finally, per constraint (5), we define the impact of the inhibitory units on the excitatory units as comprising both a subtractive and a divisive component:

$$\mathbf{h}_\ell = f(\mathbf{z}_\ell) \quad \mathbf{z}_\ell = \frac{\mathbf{g}_\ell}{\gamma_\ell} \odot (\mathbf{z}_\ell^E - \mathbf{W}_\ell^{EI} \mathbf{h}_\ell^I) + \beta_\ell \quad (1)$$

$$\text{where } \mathbf{z}_\ell^E = \mathbf{W}_\ell^{EE} \mathbf{h}_{\ell-1} \quad \mathbf{h}_\ell^I = f^I(\mathbf{z}_\ell^I) = f^I(\mathbf{W}_\ell^{IE} \mathbf{h}_{\ell-1})$$

$$\gamma_\ell = \mathbf{W}_\ell^{EI} (e^{\alpha_\ell} \odot \mathbf{h}_\ell^I)$$

where for each layer ℓ , $\beta_\ell \in \mathbb{R}^{n_e}$ is a bias, $\mathbf{g}_\ell \in \mathbb{R}_+^{n_e}$ controls the gain, γ_ℓ is the divisive inhibitory term, and $\alpha_\ell \in \mathbb{R}^{n_i}$ is a parameter that controls the strength of this divisive inhibition. In the rest of this manuscript we set f to be the rectified linear function (ReLU). Though a ReLU function is not a perfect match to the input-output properties of real neurons, it captures the essential rectification operation performed by neurons in physiologically realistic low activity regimes (Salinas & Sejnowski, 2000). In this paper, we model the inhibitory units as linear (i.e. $f^I(\mathbf{z}^I) = \mathbf{z}^I$) for mathematical

simplicity, however there is some biological justification, as the resting membrane potential of the class of fast-spiking interneurons most related to our model is relatively depolarised and their spike outputs can follow single inputs one-to-one (Hu et al., 2014; Galarreta & Hestrin, 2001).

3 PARAMETER INITIALISATION FOR DALE’S ANNS

In biology excitation and inhibition are balanced (Isaacson & Scanziani, 2011), and we use this biological property to derive appropriate weight initialisation for DANNs. First we initialise excitatory parameters using $\mathbf{W}^{\text{EE}} \stackrel{iid}{\sim} \text{Exp}(\lambda^{\text{E}})$, and then inhibitory parameters are initialised such that excitation and subtractive inhibition are balanced, i.e. $\mathbb{E}[z_k^{\text{E}}] = \mathbb{E}[(\mathbf{W}^{\text{EI}} \mathbf{z}^{\text{I}})_k]$, $\forall k$. This can be achieved in a number of ways (see appendix B.2). In line with biology, we choose to treat excitatory weights onto inhibitory and excitatory units the same, and sample $\mathbf{W}^{\text{IE}} \stackrel{iid}{\sim} \text{Exp}(\lambda^{\text{E}})$ and set $\mathbf{W}^{\text{EI}} \leftarrow \mathbf{1}/n_i$. We note that for a DANN layer with a single inhibitory neuron, e.g. at an output layer with 10 excitatory neurons, the noise inherent in sampling a single weight vector may result in a poor match between the excitatory and inhibitory inputs, so in this case we initialise \mathbf{W}^{IE} as $\frac{1}{n_e} \sum_{j=1}^{n_e} \mathbf{w}_{j,:}^{\text{EE}}$ explicitly (where $\mathbf{w}_{j,:}^{\text{EE}}$ is the j^{th} row of \mathbf{W}^{EE}).

Next, we consider the relationship between this initialisation approach and normalisation schemes (Ba et al., 2016; Ioffe & Szegedy, 2015). Normalisation acts to both center and scale the unit activities in a network such that they have mean zero and variance one. The weight initialisation given above will produce centered activities at the start of training. We can also draw a connection between the divisive inhibition and standardisation if we assume that the elements of \mathbf{x} are sampled from a rectified normal distribution, $\mathbf{x} \stackrel{iid}{\sim} \max(0, \mathcal{N}(0, \sigma_{\ell-1}^2))$. Under this assumption, the mean and variance of the excitatory input are proportional (see Appendix C). For example, if we consider the relationship $c \cdot \mathbb{E}[z_k^{\text{E}}] = \text{Var}(z_k^{\text{E}})^{1/2}$ for each unit k , we get the scalar proportionality constant $c = \sqrt{2\pi - 1}/\sqrt{d}$, as:

$$\begin{aligned} \mathbb{E}[z_k^{\text{E}}] &= d \cdot \mathbb{E}[w^{\text{EE}}] \mathbb{E}[x] \\ &= d \cdot \mathbb{E}[w^{\text{EE}}] \frac{\sigma_{\ell-1}}{\sqrt{2\pi}} \\ \text{Var}(z_k^{\text{E}}) &= d \cdot \text{Var}(w^{\text{EE}}) (\mathbb{E}[x^2] + \text{Var}(x)) \\ &= d \cdot \text{Var}(w^{\text{EE}}) \sigma_{\ell-1}^2 \frac{2\pi - 1}{2\pi} \end{aligned} \tag{2}$$

with expectation over the data and the parameters, and where w^{EE}, x refer to any element of $\mathbf{W}^{\text{EE}}, \mathbf{x}$. Therefore, since $\mathbb{E}[W^{\text{EE}}]^2 = \text{Var}(W^{\text{EE}})$ for weights drawn from an exponential distribution, we have

$$c = \frac{\text{Var}(z_k^{\text{E}})^{\frac{1}{2}}}{\mathbb{E}[z_k^{\text{E}}]} = \frac{\sqrt{2\pi - 1}}{\sqrt{d}} \tag{3}$$

This proportionality means that you can perform centering and standardisation operations using the same neurons. For DANNs, e^{α} will dictate the expected variance of the layer’s activation \mathbf{z} , as it controls the proportionality between subtractive and divisive inhibition for each inhibitory unit. If e^{α} is set to c , then the divisive inhibition approximates dividing \mathbf{z}^{E} by its standard deviation, as $\mathbb{E}[z_k^{\text{E}}] \cdot c = \mathbb{E}[\mathbf{w}_{k,:}^{\text{EI}} (e^{\alpha} \odot \mathbf{z}_k^{\text{I}})] = \mathbb{E}[\gamma_k]$. We note that due to the proportionality between \mathbf{z}^{E} ’s mean and variance, other values of e^{α} will also control the layer’s variance with depth. However, given these considerations, we initialise $e^{\alpha} \leftarrow \sqrt{2\pi - 1}/\sqrt{d}$, thereby achieving standardisation at initialisation. We find that these initialisation schemes enable DANNs to learn well. We next turn to the question of how to perform parameter updates in DANNs in order to learn well.

4 PARAMETER UPDATES FOR DALE’S ANNS

Unlike a column constrained network, a layer in a DANN is not restricted in its potential function space. This is because excitatory inputs to a layer can still have an inhibitory impact via feedforward inhibition. However, the inhibitory interneuron architecture of DANN layers introduces disparities in the degree to which updates to different parameters affect the layer’s output distribution. This can

be seen intuitively, for example if a single element of \mathbf{W}^{IE} is updated, this has an effect on each element of \mathbf{z} . Similarly, an update to w_{ij}^{EI} will change z_i depending on the alignment of x and all of the j^{th} inhibitory unit's weights. Therefore, instead of using the euclidean metric to measure distance between parameter settings, we employ an alternative approach. Similar to natural gradient methods, we use an approximation of the KL divergence of the layer's output distribution for our metric. In order to help ensure that both excitatory and inhibitory parameter updates have similar impacts on the KL divergence, we scale the updates using correction terms derived below. We provide an extended derivation of these scaling factors in the Appendix D.

Given a probability distribution parameterized by some vector θ , a second order approximation to the KL divergence for a change to the parameters θ is

$$D_{\text{KL}}[P(\mathbf{y}|\mathbf{x};\theta) \| P(\mathbf{y}|\mathbf{x};\theta + \delta)] \approx \frac{1}{2} \delta^T F(\theta) \delta \quad (4)$$

$$F(\theta) = \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x}), \mathbf{y} \sim P(\mathbf{y}|\mathbf{x};\theta)} \left[\frac{\partial \log P(\mathbf{y}|\mathbf{x};\theta)}{\partial \theta} \frac{\partial \log P(\mathbf{y}|\mathbf{x};\theta)}{\partial \theta}^T \right] \quad (5)$$

Where $F(\theta)$ is the Fisher Information matrix (or just the Fisher). In order to calculate the Fisher for the parameters of a neural network, we must interpret the network's outputs in a probabilistic manner. One approach is to view a layer's activation as parameterising a conditional distribution from the exponential family $P(\mathbf{y}|\mathbf{x};\theta) = P(\mathbf{y}|\mathbf{z})$, independent in each coordinate of $\mathbf{y}|\mathbf{z}$ (similar to a GLM, and as done in Ba et al. (2016)). The log likelihood of such a distribution can be written as¹:

$$\log P(\mathbf{y}|\mathbf{x};\theta) = \frac{\mathbf{y} \cdot \mathbf{z} - \eta(\mathbf{z})}{\phi} + c(\mathbf{y}, \phi) \quad (6)$$

$$\mathbb{E}[\mathbf{y}|\mathbf{x};\theta] = f(\mathbf{z}) = \eta'(\mathbf{z}) \quad \text{Cov}(\mathbf{y}|\mathbf{x};\theta) = \text{diag}(\phi f'(\mathbf{z})) \quad (7)$$

where $f(\mathbf{z})$ is the activation function of the layer, and ϕ, η, c define the particular distribution in the exponential family. Note that here we are taking $\eta'(\mathbf{z})$ and $f'(\mathbf{z})$ to denote $\frac{\partial \eta}{\partial \mathbf{z}}$ and $\frac{\partial f}{\partial \mathbf{z}}$, respectively.

In our networks, we have used softmax activation functions at the output and ReLU activation functions in the hidden layers. In this setting, the log likelihood of the output softmax probability layer would only be defined for a one-hot vector \mathbf{y} and would correspond to $\phi = 1$, $c(\mathbf{y}, \phi) = 0$, and $\eta(\mathbf{z}) = \log(\sum_i e^{z_i})$. For the ReLU activation functions, the probabilistic model corresponds to a tobit regression model, in which \mathbf{y} is a censored observation of a latent variable $\hat{\mathbf{y}} \sim \mathcal{N}(\mathbf{z}, \text{diag}(\phi f'(\mathbf{z})))$. In this case, one could consider either the censored or pre-censored random variable, depending on modelling preference. As it fits well with the above framework we analyze the pre-censored random variable, i.e. $f(\mathbf{z}) = \mathbf{z}$ in equation 6. Returning to the general case, the Fisher of a layer is therefore:

$$F(\theta) = \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x}), \mathbf{y} \sim P(\mathbf{y}|\mathbf{x};\theta)} \left[\frac{\partial \mathbf{z}}{\partial \theta} \frac{(\mathbf{y} - \eta'(\mathbf{z}))}{\phi} \frac{(\mathbf{y} - \eta'(\mathbf{z}))^T}{\phi} \frac{\partial \mathbf{z}^T}{\partial \theta} \right] \quad (8)$$

$$= \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[\frac{\partial \mathbf{z}}{\partial \theta} \frac{\text{diag}(f'(\mathbf{z}))}{\phi} \frac{\partial \mathbf{z}^T}{\partial \theta} \right] \quad (9)$$

To analyse the approximate KL divergence resulting from the simple case of perturbing individual parameters of a single-layer DANN, we only need to consider the diagonal entries of the Fisher:

$$D_{\text{KL}}[P_{\theta} \| P_{\theta + \delta_{\bar{\theta}}}] \approx \frac{\delta^2}{2\phi} \sum_k^{n_e} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_k) \left(\frac{\partial z_k}{\partial \theta} \right)^2 \right] \quad (10)$$

where $\delta_{\bar{\theta}}$ represents a 1-hot vector corresponding to $\bar{\theta}$, multiplied by a scalar δ . We now consider the approximate KL divergence after updates to a single element of \mathbf{W}^{EE} , \mathbf{W}^{IE} , \mathbf{W}^{EI} and α :

¹Note the general form of the exponential families is $\log P(\mathbf{y}|\mathbf{z}) = \frac{\mathbf{z} \cdot T(\mathbf{y}) - \eta(\mathbf{z})}{\phi} + c(\mathbf{y}, \phi)$ but here we only need to consider distributions where $T(y) = y$.

$$D_{\text{KL}}[P_{\theta} \parallel P_{\theta+\delta \mathbf{w}_{ij}^{\text{EE}}}] \approx \frac{\delta^2}{2\phi} \mathbb{E} \left[f'(z_i) \left(\frac{g_i}{\gamma_i} x_j \right)^2 \right] \quad (11)$$

$$D_{\text{KL}}[P_{\theta} \parallel P_{\theta+\delta \mathbf{w}_{ij}^{\text{IE}}}] \approx \frac{\delta^2}{2\phi} \sum_k^{n_e} \mathbb{E} \left[f'(z_k) \left(\frac{g_k}{\gamma_k} x_j \right)^2 (w_{ki}^{\text{EI}} a_{ki})^2 \right] \quad (12)$$

$$D_{\text{KL}}[P_{\theta} \parallel P_{\theta+\delta \mathbf{w}_{ij}^{\text{EI}}}] \approx \frac{\delta^2}{2\phi} \sum_n^d \mathbb{E} \left[f'(z_i) \left(\frac{g_i}{\gamma_i} x_n \right)^2 (w_{jn}^{\text{IE}} a_{ij})^2 \right] \\ + \frac{\delta^2}{2\phi} \sum_{n \neq m}^d \mathbb{E} \left[f'(z_i) \left(\frac{g_i}{\gamma_i} \right)^2 x_n x_m w_{jn}^{\text{IE}} w_{jm}^{\text{IE}} (a_{ij})^2 \right] \quad (13)$$

$$D_{\text{KL}}[P_{\theta} \parallel P_{\theta+\delta \alpha_i}] = \frac{\delta^2}{2\phi} \sum_k^{n_e} \sum_j^d \mathbb{E} \left[f'(z_k) \left(\frac{g_k}{\gamma_k} x_j \right)^2 w_{ki}^{\text{EI}} w_{ij}^{\text{IE}} (a_{ki} - 1)^2 \right] \\ + \frac{\delta^2}{2\phi} \sum_k^{n_e} \sum_{n \neq m}^d \mathbb{E} \left[f'(z_k) \left(\frac{g_k}{\gamma_k} \right)^2 x_n x_m w_{in}^{\text{IE}} w_{im}^{\text{IE}} (w_{ki}^{\text{EI}})^2 (a_{ki} - 1)^2 \right] \quad (14)$$

Where $a_{kj} = \frac{e^{\alpha_j}}{\gamma_k} (z_k^{\text{E}} - (\mathbf{W}^{\text{EI}} z^{\text{I}})_k) + 1$, and expectations are over the data, $\mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})}$.

Therefore, as a result of the feed-forward inhibitory architecture of DANNs, for a parameter update δ , the effect on the model's distribution will be different depending on the updated parameter-type. While the exact effect depends on the degree of co-variance between terms, the most prevalent differences between and within the excitatory and inhibitory parameters are the sums over layer input and output dimensions. For example, an inhibitory weight update of δ to w_{ij}^{IE} is expected to change the model distribution approximately n_e times more than an excitatory weight update of δ to w_{ij}^{EE} . In order to balance the impact of updating different parameter-types, we update DANN parameters after correcting for these terms: updates to \mathbf{W}^{IE} were scaled by $\sqrt{n_e}^{-1}$, \mathbf{W}^{EI} by d^{-1} and α by $(d\sqrt{n_e})^{-1}$. As a result, inhibitory unit parameters updates are scaled down relative to excitatory parameter updates. This is intriguing given the differences between inhibitory and excitatory neuron plasticity, which has been the subject of discussion in neuroscience (Kullmann & Lamsa, 2007; Kullmann et al., 2012), including the relative extent of weight changes in excitatory and inhibitory neurons (McBain et al., 1999).

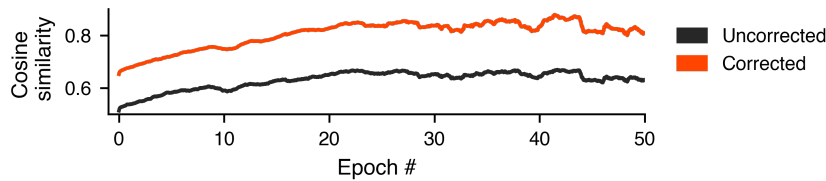


Figure 2: Empirical verification of update correction terms: Cosine of the angle between gradients multiplied by an approximation of the diagonal of the Fisher inverse for each layer, and either uncorrected gradients (black) or corrected gradients (orange) over 50 epochs. Plot displays a moving average over 500 updates

Next, we empirically verified that our heuristic correction factors captured the key differences between parameter-types in their impact on the KL divergence. To do this we compared parameter gradients before and after correction, to parameter gradients multiplied by an approximation of the diagonal of the Fisher inverse for each layer (which we refer to as Fisher corrected gradients), see Appendix E.3. The model was trained for 50 epochs on MNIST, and updated using the Fisher corrected gradients. Throughout training, we observed that the heuristic corrected gradients were more aligned to the Fisher corrected gradients than the uncorrected gradients were (Fig. 2). Thus, our derived correction factors help to balance the impact of excitatory and inhibitory updates on the network's behaviour. Below, we demonstrate that these corrections are key to getting DANNs to learn well.

5 EXPERIMENTAL RESULTS

Having derived appropriate parameter initialisation and updates for DANNs, we now explore how they compare to traditional ANNs and ColumnEi models on simple benchmark datasets. In brief, we find that column constrained models perform poorly, failing even to achieve zero training-set error, whereas DANNs perform equivalently to traditional ANNs.

5.1 IMPLEMENTATION DETAILS

All models were composed of 4 layers: in general 3 hidden layers of dimension 500 with a ReLU activation function followed by a softmax output with 10 units, and all experiments were run for 50 epochs with batch size 32. Unless stated, for DANNs and ColumnEi models, 50 inhibitory units were included per hidden layer. For DANN models, the softmax output layer was constructed with one inhibitory unit. For ColumnEi models, each hidden layer's activation is $\mathbf{z} = \mathbf{W}\mathbf{x}$ where 500 columns of \mathbf{W} were constrained to be positive and 50 negative (therefore for ColumnEi models \mathbf{h}_ℓ was of dimension 550). ColumnEi layer weights were initialised so that variance did not scale with depth and that activations were centered (see Appendix B.1 for further details). All benchmark datasets (MNIST, Kuzushiji MNIST and Fashion MNIST) were pre-processed so that pixel values were in $[0, 1]$. Learning rates were selected according to validation error averaged over 3 random seeds, after a random search (Orion; Bouthillier et al. (2019), log uniform $[10, 1e-5]$, 100 trials, 10k validation split). Selected models were then trained on test data with 6 random seeds. Plots show mean training error per epoch, and mean test set error every 200 updates over random seeds. Tables show final error mean \pm standard deviation. For further implementation details see Appendix E.

Note that because our goal in this paper is not to achieve state-of-the-art performance, we did not apply regularisation techniques, such as dropout and weight decay, or common modifications to stochastic gradient descent. Instead the goal of the experiments presented here was simply to determine whether, in the simplest test case scenario, DANNs can learn better than ColumnEi models and as well as traditional ANNs.

5.2 COMPARISON OF DANNs TO COLUMN-EI MODELS AND MLPs

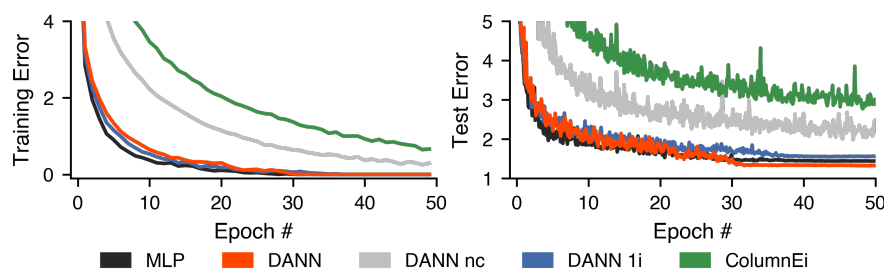


Figure 3: Model comparison on MNIST dataset. *nc* - no update corrections, *1i* - one inhibitory unit

We first compared model performance on the MNIST dataset (Fig 3). We observed that ColumnEi models generalised poorly, and failed to achieve 0 % training error within the 50 epochs. This confirms the fact that such models cannot learn as well as traditional ANNs. In contrast, we observed that DANNs performed equivalently to MLPs, and even generalised marginally better. In addition, performance was only slightly worse when DANNs were constructed with one inhibitory unit per layer. We also found that not correcting parameter updates using the corrections derived from the Fisher significantly impaired optimization, further verifying the correction factors (see also Fig 2).

Next, we compared DANN performance to MLPs trained with batch and layer normalization on more challenging benchmark datasets (Fig 4). Again we found that DANNs performed equivalently to these standard architectures, whereas ColumnEi models struggled to achieve acceptable performance.

Finally, we explored methods for improving DANN performance. First, in order to maintain the positive DANN weight constraint, if after a parameter update a weight was negative, we reset it to

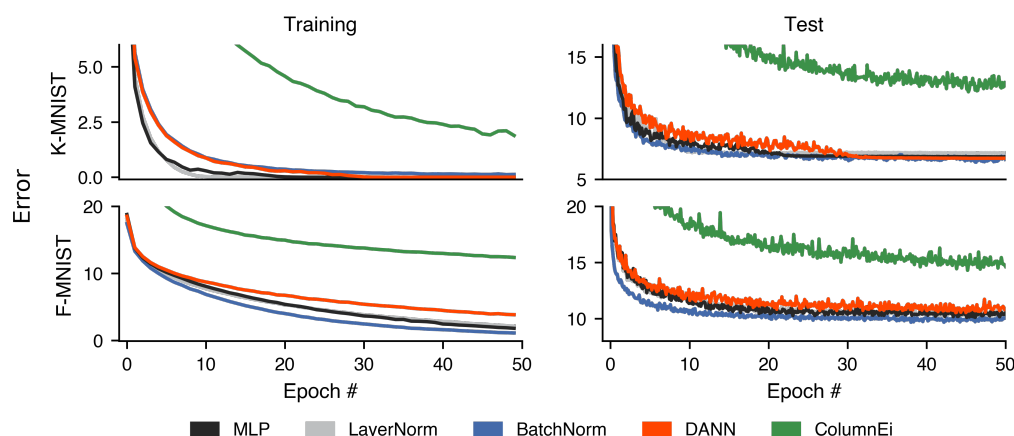


Figure 4: Model comparison on Fashion MNIST and Kuzushiji MNIST datasets.

zero, i.e. $\theta \leftarrow \max(0, \theta)$, and as a result the actual update is no longer that suggested by stochastic gradient descent. We therefore experimented with temporarily reducing the learning rate whenever this parameter clipping would reduce the cosine of the angle made between the gradient and actual updates below a certain constraint (see Appendix E.4). Second, we note that the divisive inhibition term, γ , appears in the denominator of the weight gradients (Appendix D.2) and, therefore, if γ becomes small, the gradients will become large, potentially resulting in inappropriate parameter updates. We therefore wondered if gradient normalisation would be particularly effective for DANNs. We tested both of these modifications to DANNs trained on FashionMNIST (Supp. Fig. 5). However, we found that they provided no observable improvement, indicating that the loss landscape and gradients were well behaved over optimization.

Altogether, our results demonstrate that: (1) the obvious approach to creating ANNs that obey Dale’s principle (ColumnEi models) do not learn as well as traditional ANNs, (2) DANNs learn better than ColumnEi models and as well as traditional ANNs, (3) DANN learning is significantly improved by taking appropriate steps to scale updates in excitatory and inhibitory units appropriately.

6 DISCUSSION

Here we presented DANNs, a novel ANN architecture with separate inhibitory and excitatory units. We derived appropriate parameter initialisation and update rules and showed experimentally that, unlike ANNs where some columns are simply constrained to be positive or negative, DANNs perform equivalently to traditional ANNs on benchmark datasets. These results are important as they are, as far as we know, the first example of an ANN architecture that fully adheres to Dale’s law *without sacrificing learning performance*.

There are a number of avenues for future work building upon DANNs, the most obvious of which are to extend and incorporate DANNs into recurrent and convolution neural networks. An alternative direction is explore more fully the balance of excitation and inhibition during optimization, as while DANNs are initialised such that these are balanced, the inhibitory parameters are updated during training, and the model is free to diverge from this initialisation. As a result, the distribution of layer activations may be unstable over successive parameter updates, potentially harming optimization. In the brain, a variety of homeostatic plasticity mechanisms stabilize neuronal activity. For example, reducing excitatory input naturally results in a reduction in inhibition in real neural circuits (Tien & Kerschensteiner, 2018). It would therefore be interesting to test the inclusion of a homeostatic loss to encourage inhibition to track excitation throughout training.

In summary, DANNs sit at the intersection of a number of programs of research. First, they are a new architecture that obeys Dale’s principle, but which can still learn well, allowing researchers to more directly compare trained ANNs to real neural data (Schrimpf et al., 2018; Yamins et al., 2014). Second, DANNs contribute towards computational neuroscience and machine learning work on inhibitory interneurons in ANNs, and in general towards the role of inhibitory circuits and plasticity in neural

computation (Song et al., 2016; Sacramento et al., 2018; Costa et al., 2017; Payeur et al., 2020; Atallah et al., 2012; Barron et al., 2017). Finally, the inhibition in DANNs also has an interesting connection to normalisation methods used to improving learning in deep networks (Ioffe & Szegedy, 2015; Wu & He, 2018; Ba et al., 2016). As DANNs tie these distinct programs of research together into a single model, we hope they can serve as a basis for future research at the intersection of deep learning and neuroscience.

REFERENCES

- Daniel J Amit, C Campbell, and KYM Wong. The interaction space of neural networks with sign-constrained synapses. *Journal of Physics A: Mathematical and General*, 22(21):4687, 1989.
- Bassam V Atallah, William Bruns, Matteo Carandini, and Massimo Scanziani. Parvalbumin-expressing interneurons linearly transform cortical responses to visual stimuli. *Neuron*, 73(1): 159–170, 2012.
- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Helen C Barron, Tim P Vogels, Timothy E Behrens, and Mani Ramaswami. Inhibitory engrams in perception and memory. *Proceedings of the National Academy of Sciences*, 114(26):6666–6674, 2017.
- Sergey Bartunov, Adam Santoro, Blake Richards, Luke Marris, Geoffrey E Hinton, and Timothy Lillcrap. Assessing the scalability of biologically-motivated deep learning algorithms and architectures. In *Advances in Neural Information Processing Systems*, pp. 9368–9378, 2018.
- Xavier Bouthillier, Christos Tsirigotis, François Corneau-Tremblay, Pierre Delaunay, Reyhane Askari, Dendi Suhubdy, Michael Noukhovitch, Dmitriy Serdyuk, Arnaud Bergeron, Peter Henderson, Pascal Lamblin, Mirko Bronzi, and Christopher Beckham. Orion - asynchronous distributed hyperparameter optimization, October 2019. URL <https://doi.org/10.5281/zenodo.3478592>.
- Rui Costa, Ioannis Alexandros Assael, Brendan Shillingford, Nando de Freitas, and Tim Vogels. Cortical microcircuits as gated-recurrent neural networks. In *Advances in neural information processing systems*, pp. 272–283, 2017.
- Peter W Donhauser and Sylvain Baillet. Two distinct neural timescales for predictive speech processing. *Neuron*, 105(2):385–393, 2020.
- 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.
- Mario Galarreta and Shaul Hestrin. Spike transmission and synchrony detection in networks of gabaergic interneurons. *Science*, 292(5525):2295–2299, 2001.
- Hua Hu, Jian Gan, and Peter Jonas. Fast-spiking, parvalbumin+ gabaergic interneurons: From cellular design to microcircuit function. *Science*, 345(6196):1255263, 2014.
- Alessandro Ingrosso and LF Abbott. Training dynamically balanced excitatory-inhibitory networks. *PloS one*, 14(8):e0220547, 2019.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Jeffrey S Isaacson and Massimo Scanziani. How inhibition shapes cortical activity. *Neuron*, 72(2): 231–243, 2011.
- Alexander JE Kell, Daniel LK Yamins, Erica N Shook, Sam V Norman-Haignere, and Josh H McDermott. A task-optimized neural network replicates human auditory behavior, predicts brain responses, and reveals a cortical processing hierarchy. *Neuron*, 98(3):630–644, 2018.

- Tim Christian Kietzmann, Patrick McClure, and Nikolaus Kriegeskorte. Deep neural networks in computational neuroscience. *BioRxiv*, pp. 133504, 2018.
- Jonas Kubilius, Martin Schrimpf, Kohitij Kar, Rishi Rajalingham, Ha Hong, Najib Majaj, Elias Issa, Pouya Bashivan, Jonathan Prescott-Roy, Kailyn Schmidt, et al. Brain-like object recognition with high-performing shallow recurrent anns. In *Advances in Neural Information Processing Systems*, pp. 12805–12816, 2019.
- Dimitri M Kullmann and Karri P Lamsa. Long-term synaptic plasticity in hippocampal interneurons. *Nature Reviews Neuroscience*, 8(9):687–699, 2007.
- Dimitri M Kullmann, Alexandre W Moreau, Yamina Bakiri, and Elizabeth Nicholson. Plasticity of inhibition. *Neuron*, 75(6):951–962, 2012.
- Timothy P Lillicrap, Adam Santoro, Luke Marris, Colin J Akerman, and Geoffrey Hinton. Backpropagation and the brain. *Nature Reviews Neuroscience*, pp. 1–12, 2020.
- Joana Lourenço, Angela Michela De Stasi, Charlotte Deleuze, Mathilde Bigot, Antonio Pazienti, Andrea Aguirre, Michele Giugliano, Srdjan Ostojic, and Alberto Bacci. Modulation of coordinated activity across cortical layers by plasticity of inhibitory synapses. *Cell reports*, 30(3):630–641, 2020.
- James Martens. New insights and perspectives on the natural gradient method. *arXiv preprint arXiv:1412.1193*, 2014.
- James Martens and Roger Grosse. Optimizing neural networks with kronecker-factored approximate curvature. In *International conference on machine learning*, pp. 2408–2417, 2015.
- Chris J McBain, Tamas F Freund, and Istvan Mody. Glutamatergic synapses onto hippocampal interneurons: precision timing without lasting plasticity. *Trends in neurosciences*, 22(5):228–235, 1999.
- Jonathan A Michaels, Stefan Schaffelhofer, Andres Agudelo-Toro, and Hansjörg Scherberger. A neural network model of flexible grasp movement generation. *bioRxiv*, pp. 742189, 2019.
- Thomas Miconi. Biologically plausible learning in recurrent neural networks reproduces neural dynamics observed during cognitive tasks. *Elife*, 6:e20899, 2017.
- Sun Minni, Li Ji-An, Theodore Moskovitz, Grace Lindsay, Kenneth Miller, Mario Dipoppa, and Guangyu Robert Yang. Understanding the functional and structural differences across excitatory and inhibitory neurons. 2019.
- Aran Nayeibi, Daniel Bear, Jonas Kubilius, Kohitij Kar, Surya Ganguli, David Sussillo, James J DiCarlo, and Daniel L Yamins. Task-driven convolutional recurrent models of the visual system. In *Advances in Neural Information Processing Systems*, pp. 5290–5301, 2018.
- Christopher Parisien, Charles H Anderson, and Chris Eliasmith. Solving the problem of negative synaptic weights in cortical models. *Neural computation*, 20(6):1473–1494, 2008.
- Alexandre Payeur, Jordan Guerguiev, Friedemann Zenke, Blake Richards, and Richard Naud. Burst-dependent synaptic plasticity can coordinate learning in hierarchical circuits. *bioRxiv*, 2020.
- Frédéric Pouille, Antonia Marin-Burgin, Hillel Adesnik, Bassam V Atallah, and Massimo Scanziani. Input normalization by global feedforward inhibition expands cortical dynamic range. *Nature neuroscience*, 12(12):1577, 2009.
- Frederic Pouille, Oliver Watkinson, Massimo Scanziani, and Andrew J Trevelyan. The contribution of synaptic location to inhibitory gain control in pyramidal cells. *Physiological reports*, 1(5), 2013.
- Blake A Richards, Timothy P Lillicrap, Philippe Beaudoin, Yoshua Bengio, Rafal Bogacz, Amelia Christensen, Claudia Clopath, Rui Ponte Costa, Archy de Berker, Surya Ganguli, et al. A deep learning framework for neuroscience. *Nature neuroscience*, 22(11):1761–1770, 2019.

- João Sacramento, Rui Ponte Costa, Yoshua Bengio, and Walter Senn. Dendritic cortical microcircuits approximate the backpropagation algorithm. In *Advances in Neural Information Processing Systems*, pp. 8721–8732, 2018.
- Emilio Salinas and Terrence J Sejnowski. Impact of correlated synaptic input on output firing rate and variability in simple neuronal models. *Journal of neuroscience*, 20(16):6193–6209, 2000.
- Martin Schrimpf, Jonas Kubilius, Ha Hong, Najib J Majaj, Rishi Rajalingham, Elias B Issa, Kohitij Kar, Pouya Bashivan, Jonathan Prescott-Roy, Kailyn Schmidt, et al. Brain-score: Which artificial neural network for object recognition is most brain-like? *BioRxiv*, pp. 407007, 2018.
- Bryan A Seybold, Elizabeth AK Phillips, Christoph E Schreiner, and Andrea R Hasenstaub. Inhibitory actions unified by network integration. *Neuron*, 87(6):1181–1192, 2015.
- H Francis Song, Guangyu R Yang, and Xiao-Jing Wang. Training excitatory-inhibitory recurrent neural networks for cognitive tasks: a simple and flexible framework. *PLoS computational biology*, 12(2):e1004792, 2016.
- Nai-Wen Tien and Daniel Kerschensteiner. Homeostatic plasticity in neural development. *Neural development*, 13(1):1–7, 2018.
- Robin Tremblay, Soohyun Lee, and Bernardo Rudy. Gabaergic interneurons in the neocortex: from cellular properties to circuits. *Neuron*, 91(2):260–292, 2016.
- James CR Whittington and Rafal Bogacz. Theories of error back-propagation in the brain. *Trends in cognitive sciences*, 2019.
- Nathan R Wilson, Caroline A Runyan, Forea L Wang, and Mriganka Sur. Division and subtraction by distinct cortical inhibitory networks in vivo. *Nature*, 488(7411):343–348, 2012.
- Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 3–19, 2018.
- Daniel LK Yamins and James J DiCarlo. Using goal-driven deep learning models to understand sensory cortex. *Nature neuroscience*, 19(3):356, 2016.
- Daniel LK Yamins, Ha Hong, Charles F Cadieu, Ethan A Solomon, Darren Seibert, and James J DiCarlo. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences*, 111(23):8619–8624, 2014.

SUPPLEMENTARY MATERIAL

A SUPPLEMENTARY RESULTS

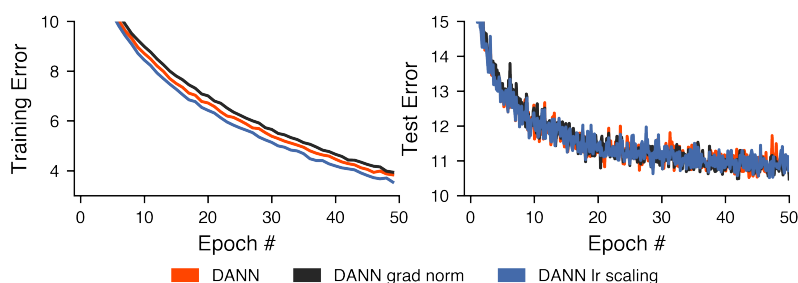


Figure 5: DANNs trained on Fashion MNIST with gradient normalisation and learning rate scaling.

Table 1: MNIST results

Model	# inhib	Learning rate	Train error	Test error
MLP	0	0.2976	0.0 ± 0.0	1.44 ± 0.034
DANN	50	0.1107	0.0 ± 0.0	1.325 ± 0.066
DANN nc	50	0.003981	0.293 ± 0.039	2.167 ± 0.113
DANN li	1	0.05327	0.0 ± 0.0	1.56 ± 0.041
ColumnEi	50	0.05273	0.662 ± 0.071	3.035 ± 0.231

Table 2: K-MNIST results

Model	# inhib	Learning rate	Train error	Test error
MLP	0	0.213	0.0 ± 0.0	6.842 ± 0.119
LayerNorm	0	0.1166	0.0 ± 0.0	7.171 ± 0.134
BatchNorm	0	0.7235	0.123 ± 0.017	6.747 ± 0.129
DANN	50	0.156	0.0 ± 0.0	6.728 ± 0.23
ColumnEi	50	0.05426	1.886 ± 0.191	12.763 ± 0.44

Table 3: Fashion MNIST results

Model	# inhib	Learning rate	Train error	Test error
MLP	0	0.08403	1.793 ± 0.119	10.626 ± 0.466
LayerNorm	0	0.04743	2.136 ± 0.076	10.445 ± 0.455
BatchNorm	0	0.1098	1.104 ± 0.044	9.992 ± 0.218
DANN	50	0.01973	3.832 ± 0.031	10.962 ± 0.365
ColumnEi	50	0.02265	12.365 ± 0.217	14.986 ± 0.674

B PARAMETER INITIALISATIONS

In this section we provide further details regarding parameter initialisations.

Throughout we assume that the elements of the input, \mathbf{x} , to a layer ℓ are *iid* and distributed $\mathcal{N}(0, \sigma_{\ell-1}^2)$. Therefore x will follow a rectified normal distribution:

$$\begin{aligned}\mathbb{E}[x] &= \int_0^\infty x \cdot \frac{e^{-x^2/2\sigma_{\ell-1}^2}}{\sigma_{\ell-1}\sqrt{2\pi}} dx = \frac{\sigma_{\ell-1}}{\sqrt{2\pi}} \\ \mathbb{E}[x^2] &= \int_0^\infty x^2 \cdot \frac{e^{-x^2/2\sigma_{\ell-1}^2}}{\sigma_{\ell-1}\sqrt{2\pi}} dx = \frac{\sigma_{\ell-1}^2}{2} \\ \text{Var}(x) &= \mathbb{E}[x^2] - \mathbb{E}[x]^2 = \sigma_{\ell-1}^2 \frac{\pi - 1}{2\pi} \\ \text{Var}(x) + \mathbb{E}[x]^2 &= \sigma_{\ell-1}^2 \frac{2\pi - 1}{2\pi}\end{aligned}\tag{15}$$

where here, and throughout the text, non-indexed non-bold to refers to any element of a vector or matrix, e.g $\mathbb{E}[x]$ refers to the expectation of any element of \mathbf{x} , $\text{Var}(x)$ refers to the variance of any element of \mathbf{x} , etc. In addition, for all models we draw positively constrained weights *iid* from exponential distributions, and make use of the following properties for $w \sim \text{Exp}(\lambda)$

$$\begin{aligned}\mathbb{E}[w] &= \frac{1}{\lambda} \\ \text{Var}(w) &= \frac{1}{\lambda^2} = \mathbb{E}[w]^2 \\ \mathbb{E}[w^2] &= \frac{2!}{\lambda^2} = 2\text{Var}(w) = 2\mathbb{E}[w]^2\end{aligned}\tag{16}$$

B.1 COLUMN CONSTRAINED EI MODELS AND WEIGHT INITIALISATION

Here we provide detail on the parameter initialisation of column constrained models. Layer activations are $\mathbf{z} = \mathbf{W}\mathbf{x}$ where columns of \mathbf{W} are constrained to be positive or negative. Therefore, for convenience, let us denote $\mathbf{W} = [\mathbf{W}^+, \mathbf{W}^-]$, and $\mathbf{x} = [\mathbf{x}^E, \mathbf{x}^I]$, and we assume $\mathbf{x}_i^E, \mathbf{x}_j^I$ are *iid* $\forall i, j$. Note for this model, $n_e + n_i = d$, the input dimensionality. As for DANN models, throughout training we preserve the sign constraints of the weights by resetting weights using rectification around zero, i.e. $\mathbf{W}^+ \leftarrow \max(0, \mathbf{W}^+)$, $\mathbf{W}^- \leftarrow \min(0, \mathbf{W}^-)$.

At initialisation for the column constrained model for each layer we require $\mathbb{E}[z_k] = 0$, $\text{Var}(z_k) = \sigma_{\ell-1}^2$.

$$\begin{aligned}\mathbb{E}[z_k] &= n_e \mathbb{E}[w^+] \mathbb{E}[x] - n_i \mathbb{E}[w^-] \mathbb{E}[x] \\ n_e \mathbb{E}[w^+] \mathbb{E}[x] &= n_i \mathbb{E}[w^-] \mathbb{E}[x] \\ \mathbb{E}[w^-] &= \mathbb{E}[w^+] \frac{n_e}{n_i}\end{aligned}\tag{17}$$

Where w^+, w^- refer to any element of $\mathbf{W}^+, \mathbf{W}^-$.

$$\begin{aligned}\text{Var}(z_k) &= \sum_i^{n_e} \text{Var}(w_{ki}^+ x_i^E) + \sum_j^{n_i} \text{Var}(w_{kj}^- x_j^I) \\ &= n_e \text{Var}(w^+ x) + n_i \text{Var}(w^- x) \\ &= n_e (\mathbb{E}[w^+]^2 \text{Var}(x) + \text{Var}(w^+) \mathbb{E}[x]^2 + \text{Var}(w^+) \text{Var}(x)) \\ &\quad + n_i (\mathbb{E}[w^-]^2 \text{Var}(x) + \text{Var}(w^-) \mathbb{E}[x]^2 + \text{Var}(w^-) \text{Var}(x))\end{aligned}\tag{18}$$

As weights are drawn from an exponential distribution, $\text{Var}(w^+) = \mathbb{E}[w^+]^2$, we have

$$\begin{aligned}\text{Var}(z_k) &= n_e \mathbb{E}[w^+]^2 (2\text{Var}(x) + \mathbb{E}[x]^2) + n_i \mathbb{E}[w^-]^2 (2\text{Var}(x) + \mathbb{E}[x]^2) \\ &= n_e \mathbb{E}[w^+]^2 (\mathbb{E}[x^2] + \text{Var}(x)) + n_i \mathbb{E}[w^-]^2 (\mathbb{E}[x^2] + \text{Var}(x)) \\ &= (\mathbb{E}[x^2] + \text{Var}(x)) (n_e \mathbb{E}[w^+]^2 + n_i \mathbb{E}[w^-]^2) \\ &= \sigma_{\ell-1}^2 \left(\frac{2\pi-1}{2\pi} \right) \mathbb{E}[w^+]^2 \left(n_e + \frac{n_e^2}{n_i} \right)\end{aligned}\tag{19}$$

Therefore $\mathbb{E}[w^+] = 1 / \left(\frac{2\pi-1}{2\pi} \right) \left(n_e + \frac{n_e^2}{n_i} \right)$

Note that as the input to the network is all positive, the first weight matrix has no negative columns. We therefore use the bias vector to center the activations of the first layer (in other layers it is initialised to zeros).

$$\mathbb{E}[z_k] = n_e \mathbb{E}[w^+] \mathbb{E}[x] + \beta_k\tag{20}$$

Therefore we initialise all elements of β to $-n_e \mathbb{E}[w^+] \mathbb{E}[x]$

B.2 INITIALISATION OF DANN INHIBITORY WEIGHTS FOR BALANCED EXCITATION AND SUBTRACTIVE INHIBITION

Here provide details of inhibitory parameter initialisation such that $\mathbb{E}[z_k^E] = \mathbb{E}[(\mathbf{W}^{EI} \mathbf{z}^I)_k]$, for $\mathbf{W}^{EE} \stackrel{iid}{\sim} \text{Exp}(\lambda^E)$.

$$\begin{aligned}\mathbb{E}[z_k^E] &= \mathbb{E}\left[\sum_i^d w_{ki}^{EE} x_i\right] = d \frac{1}{\lambda^E} \mathbb{E}[x] \\ \mathbb{E}[(\mathbf{W}^{EI} \mathbf{z}^I)_k] &= \mathbb{E}\left[\sum_j^{n_i} w_{kj}^{EI} \sum_i^d w_{ji}^{EI} x_i\right] = n_i \mathbb{E}[w^{EI}] d \mathbb{E}[w^{IE}] \mathbb{E}[x]\end{aligned}\tag{21}$$

These expectaions are equal when both sets of excitatory weights are drawn from the same distribution, $\mathbf{W}^{IE} \stackrel{iid}{\sim} \text{Exp}(\lambda^E)$ and $\mathbf{W}^{EI} \leftarrow 1/n_i$. Or alternatively, inhibitory weights can both drawn from the same distribution, $\mathbf{W}^{IE}, \mathbf{W}^{EI} \stackrel{iid}{\sim} \text{Exp}(\sqrt{\lambda^E n_i})$. Note, that although the above always holds in expectation, in the case of a multiple inhibitory units we can apply the law of large numbers to conclude that the subtractive inhibition and excitatory input will be approximately equal.

Note that while this initialisation is general to different settings of λ^E , we initialise $\lambda^E \leftarrow \sqrt{d(2\pi-1)}/\sqrt{2\pi}$ (see section C.1).

C PROPORTIONAL RELATIONSHIP BETWEEN EXCITATORY INPUT MEAN AND VARIANCE

Here we provide further details regarding the proportionality between \mathbf{z}^E 's mean and variance. This proportionality constant depends on which statistic or distribution that is of interest for activations (e.g. layer-statistics or unit batch-statistics as in layer and batch normalisation).

C.1 UNIT STATISTICS OVER DATA AND PARAMETER DISTRIBUTIONS

As discussed in the main text, if we consider $c \cdot \mathbb{E}[z_k^E] = \text{Var}(z_k^E)^{1/2}$ for a unit k , with expectation over the data and parameters, $c = \sqrt{2\pi - 1}/\sqrt{d}$:

$$\begin{aligned}
 \mathbb{E}[z_k^E] &= d \cdot \mathbb{E}[w^{\text{EE}}] \mathbb{E}[x] \\
 &= d \cdot \mathbb{E}[w^{\text{EE}}] \frac{\sigma_{\ell-1}}{\sqrt{2\pi}} \\
 \text{Var}(z_k^E) &= \text{Var}\left(\sum_i^d w_{ki}^{\text{EE}} x_i\right) \\
 &= d \cdot \text{Var}(w^{\text{EE}} x) \\
 &= d \cdot \text{Var}(w^{\text{EE}}) \mathbb{E}[x^2] + d \cdot \text{Var}(x) \mathbb{E}[w^{\text{EE}}]^2 \\
 &= d \cdot \text{Var}(w^{\text{EE}}) (\mathbb{E}[x^2] + \text{Var}(x)) \\
 &= d \cdot \text{Var}(w^{\text{EE}}) \sigma_{\ell-1}^2 \frac{2\pi - 1}{2\pi}
 \end{aligned} \tag{22}$$

Where $\mathbb{E}[w^{\text{EE}}]^2 = \text{Var}(w^{\text{EE}})$ for weights drawn from an exponential distribution. Therefore

$$\begin{aligned}
 \mathbb{E}[z_k^E] \cdot c &= \sqrt{\text{Var}(z_k^E)} \\
 c &= \frac{\sqrt{2\pi - 1}}{\sqrt{d}}
 \end{aligned} \tag{23}$$

Additionally, we see that for $\text{Var}(z_k^E) = \sigma_{\ell-1}^2$ the variance of the distribution that elements of \mathbf{W}^{EE} are drawn from should be

$$\text{Var}(w^{\text{EE}}) = \frac{2\pi}{d \cdot (2\pi - 1)} \tag{24}$$

and so we can set $\lambda^E \leftarrow \sqrt{d(2\pi - 1)}/\sqrt{2\pi}$, for $\text{Var}(z_k^E) = \sigma_{\ell-1}^2$.

C.2 UNIT STATISTICS OVER THE DATA DISTRIBUTION

If instead we consider a unit k , with excitatory weights \mathbf{w}_k^{EE} and expectation and variance taken only over the data we have the approximation:

$$\begin{aligned}
 \mathbb{E}[z_k^E] &= \mathbb{E}[x] \sum_i^d w_{ki}^{\text{EE}} \\
 &\approx d \cdot \mathbb{E}[x] \mathbb{E}[w^{\text{EE}}] \\
 &= d \cdot \frac{\sigma_{\ell-1}}{\sqrt{2\pi}} \mathbb{E}[w^{\text{EE}}]
 \end{aligned} \tag{25}$$

Likewise the variance over the data can be approximated as

$$\begin{aligned}
 \text{Var}(z_k^E) &= \text{Var}(x) \sum_i^d (w_{ki}^{\text{EE}})^2 \\
 &\approx d \cdot \text{Var}(x) \cdot \mathbb{E}[(w^{\text{EE}})^2] \\
 &= d \cdot \sigma_{\ell-1}^2 \frac{\pi - 1}{2\pi} \cdot 2 \cdot \mathbb{E}[w^{\text{EE}}]^2
 \end{aligned} \tag{26}$$

Therefore

$$\begin{aligned} \mathbb{E}_{x \sim p(x)}[z_k^E] \cdot c &= \sqrt{\text{Var}_{x \sim p(x)}[z_k^E]} \\ c &\approx \frac{\sqrt{2\pi - 2}}{\sqrt{d}} \end{aligned} \quad (27)$$

C.3 LAYER STATISTICS OVER THE DATA AND PARAMETER DISTRIBUTIONS

Alternatively we can consider the mean and variance of the layer statistics $\mu_{\mathbf{z}^E}$, $\sigma_{\mathbf{z}^E}$ as calculated if one was to apply layer normalisation to \mathbf{z}^E . Here again, these statistics are proportionally related, but with the constant $\sqrt{\pi}/\sqrt{d}$.

If we were to apply layer normalisation to \mathbf{z}^E , the layer statistics would be as follows:

$$\mathbf{z} = \frac{\mathbf{g}}{\sigma_{\mathbf{z}^E}}(\mathbf{z}^E - \mu_{\mathbf{z}^E}) + \beta \quad \mu_{\mathbf{z}^E} = \frac{1}{n_e} \sum_j^{n_e} z_j^E = \frac{1}{n_e} \sum_j^{n_e} \mathbf{w}_{j,:}^{\text{EE}} \mathbf{x} \quad \sigma_{\mathbf{z}^E}^2 = \frac{1}{n_e - 1} \sum_j^{n_e} (z_j^E - \mu_{\mathbf{z}^E})^2 \quad (28)$$

We now derive the relationship that the expectation of layer statistics are proportionally related by $\mathbb{E}[\mu_{\mathbf{z}^E}] \cdot \sqrt{\pi}/\sqrt{d} = \mathbb{E}(\sigma_{\mathbf{z}^E}^2)^{1/2}$. The expectation of $\mathbb{E}[\mu_{\mathbf{z}^E}]$ is straightforward:

$$\mathbb{E}[\mu_{\mathbf{z}^E}] = d \cdot \mathbb{E}[w^{\text{EE}}] \cdot \mathbb{E}[x] \quad (29)$$

Turning to the derivation of $\mathbb{E}[\sigma_{\mathbf{z}^E}^2]$:

$$\mathbb{E}[\sigma_{\mathbf{z}^E}^2] = \mathbb{E}\left[\frac{1}{n_e - 1} \sum_i^{n_e} (z_i^E - \mu_{\mathbf{z}^E})^2\right] \quad (30)$$

$$= \frac{1}{n_e - 1} \sum_i^{n_e} \mathbb{E}\left[\left(\mathbf{w}_{i,:}^{\text{EE}} \mathbf{x} - \frac{1}{n_e} \sum_j^{n_e} \mathbf{w}_{j,:}^{\text{EE}} \mathbf{x}\right)^2\right] \quad (31)$$

$$= \frac{1}{n_e - 1} \sum_i^{n_e} \mathbb{E}[(\hat{z}_i)^2] \quad (32)$$

where we have defined $\hat{z}_i = \mathbf{w}_{i,:}^{\text{EE}} \mathbf{x} - \frac{1}{n_e} \sum_j^{n_e} \mathbf{w}_{j,:}^{\text{EE}} \mathbf{x}$. We can obtain $\mathbb{E}[(\hat{z}_i)^2]$ by deriving $\mathbb{E}[\hat{z}_i]$ and $\text{Var}(\hat{z}_i)$. As

$$\hat{w}_{ij} = w_{ij}^{\text{EE}} - \frac{1}{n_e} \sum_k^{n_e} w_{kj}^{\text{EE}} = \frac{1}{n_e} \sum_{k=1, k \neq i}^{n_e} (w_{ij}^{\text{EE}} - w_{kj}^{\text{EE}}) \quad (33)$$

we see that $\mathbb{E}[\hat{w}_{ij}] = 0$, and therefore $\mathbb{E}[\hat{z}_i] = 0$. For the variance of $\text{Var}(\hat{z}_i)$ we start with $\text{Var}(\hat{w}_{ij})$.

$$\text{Var}(\hat{w}_{ij}) = \frac{1}{n_e^2} \text{Var}\left(\sum_{k=1, k \neq i}^{n_e} (w_{ij}^{\text{EE}} - w_{kj}^{\text{EE}})\right) \quad (34)$$

$$= \frac{1}{n_e^2} \left(\sum_{k=1, k \neq i}^{n_e} \text{Var}(w_{ij}^{\text{EE}} - w_{kj}^{\text{EE}}) + \sum_{\substack{k=1, k'=1 \\ k, k' \neq i, k \neq k'}}^{n_e} \text{Cov}(w_{ij}^{\text{EE}} - w_{kj}^{\text{EE}}, w_{ij}^{\text{EE}} - w_{k'j}^{\text{EE}}) \right) \quad (35)$$

$$= \frac{1}{n_e^2} ((n_e - 1)2\text{Var}(w^{\text{EE}}) + (n_e - 1)(n_e - 2)\text{Var}(w^{\text{EE}})) \quad (36)$$

$$= \frac{(n_e - 1)}{n_e} \text{Var}(w^{\text{EE}}) \quad (37)$$

For $i \leq n_e$ we calculate $\text{Var}(\hat{z}_i)$, keeping in mind that for $i \leq n_e, j \leq d$, x_j are iid, and equation (33) shows that \hat{w}_{ij} are iid in the j 'th coordinate, so we see that

$$\text{Var}(\hat{z}_i) = \text{Var}\left(\sum_j^d \hat{w}_{ij} x_j\right) = d \text{Var}(\hat{w}x) \quad (38)$$

Remembering the values of $\mathbb{E}[\hat{w}]$, $\text{Var}(\hat{w})$, that $\mathbb{E}[X^2] = \text{Var}(X) + \mathbb{E}[X]^2$, and for independent X, Y , $\text{Var}(XY) = \text{Var}(X)\text{Var}(Y) + \text{Var}(X)\mathbb{E}[Y]^2 + \text{Var}(Y)\mathbb{E}[X]^2$, we have

$$\text{Var}(\hat{z}_i) = d(\text{Var}(\hat{w})\text{Var}(x) + \text{Var}(\hat{w})\mathbb{E}[x]^2 + \text{Var}(x)\mathbb{E}[\hat{w}]^2) \quad (39)$$

$$= d(\text{Var}(\hat{w})\mathbb{E}[x^2] + \text{Var}(x)\mathbb{E}[\hat{w}]^2) = d \text{Var}(\hat{w})\mathbb{E}[x^2] \quad (40)$$

$$= \frac{d(n_e - 1)}{n_e} \text{Var}(w^{\text{EE}})\mathbb{E}[x^2] \quad (41)$$

Now putting these terms together we can derive $\mathbb{E}[\sigma_{\mathbf{z}^{\text{E}}}^2]$.

$$\mathbb{E}[\sigma_{\mathbf{z}^{\text{E}}}^2] = \frac{1}{n_e - 1} \sum_i^{n_e} \mathbb{E}[(\hat{z}_i)^2] \quad (42)$$

$$= \frac{1}{n_e - 1} \sum_i^{n_e} \text{Var}(\hat{z}_i) \quad (43)$$

$$= d \cdot \text{Var}(w^{\text{EE}}) \cdot \mathbb{E}[x^2] \quad (44)$$

Therefore returning to $\mathbb{E}[\mu_{\mathbf{z}^{\text{E}}}] \cdot c = \mathbb{E}(\sigma_{\mathbf{z}^{\text{E}}}^2)^{1/2}$ and keeping in mind that the variance of an exponential random variable is it's mean squared,

$$c = \frac{(d \cdot \text{Var}(w^{\text{EE}})\mathbb{E}[x^2])^{1/2}}{d \cdot \mathbb{E}[w^{\text{EE}}] \cdot \mathbb{E}[x]} = \frac{\sqrt{\mathbb{E}[x^2]}}{\sqrt{d} \cdot \mathbb{E}[x]} \quad (45)$$

We have assumed that x follows a rectified normal distribution. Therefore, $\mathbb{E}[x] = \frac{\sigma_{l-1}}{\sqrt{2\pi}}$, $\mathbb{E}[x^2] = \frac{\sigma_{l-1}^2}{2}$. Resulting in:

$$c = \frac{\sqrt{\pi}}{\sqrt{d}} \quad (46)$$

We note that for a DANN layer with a single inhibitory unit, $\mu_{\mathbf{z}^{\text{E}}} = z^{\text{I}}$ as $\mathbf{W}^{\text{IE}} \leftarrow \frac{1}{n_e} \sum_j^{n_e} \mathbf{w}_{j,:}^{\text{EE}}$, and $\mathbf{W}^{\text{EI}} \leftarrow \mathbb{1}$. Therefore DANN divisive inhibition, γ , can be made equivalent to layer standard deviation at initialisation in expectation if $e^{\alpha} \leftarrow c$. However, these calculations apply for the case of multiple interneuron if one makes the approximation $\mu_{\mathbf{z}^{\text{E}}} \approx (\mathbf{W}^{\text{EI}}\mathbf{W}^{\text{IE}}\mathbf{x})_i$ for any i .

D PARAMETER UPDATES AND FISHER INFORMATION MATRIX

D.1 LAYER FISHER INFORMATION MATRIX

We view a layer's activation as parameterising a conditional distribution from the exponential family $P(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}) = P(\mathbf{y}|\mathbf{z})$, independent in each coordinate of $\mathbf{y}|\mathbf{z}$.

$$\log P(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}) = \frac{\mathbf{y} \cdot \mathbf{z} - \eta(\mathbf{z})}{\phi} + c(\mathbf{y}, \phi) \quad (47)$$

$$\mathbb{E}[\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}] = f(\mathbf{z}) = \eta'(\mathbf{z}) \quad \text{Cov}(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}) = \text{diag}(\phi f'(\mathbf{z})) \quad (48)$$

where $f(\mathbf{z})$ is the activation function of the layer, and ϕ, η, c define the particular distribution in the exponential family. Note we take $\eta'(\mathbf{z}), f'(\mathbf{z})$ to denote the $\frac{\partial \eta}{\partial \mathbf{z}}, \frac{\partial f}{\partial \mathbf{z}}$.

$F(\boldsymbol{\theta})$ is defined as:

$$F(\boldsymbol{\theta}) = \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x}), \mathbf{y} \sim P(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta})} \left[\frac{\partial \log P(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \frac{\partial \log P(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta})}{\partial \boldsymbol{\theta}}^T \right] \quad (49)$$

As

$$\begin{aligned} \frac{\partial}{\partial \boldsymbol{\theta}} \log P(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}) &= \frac{\partial \mathbf{z}}{\partial \boldsymbol{\theta}} \frac{\partial}{\partial \mathbf{z}} \left[\frac{\mathbf{y} \cdot \mathbf{z} - \eta(\mathbf{z})}{\phi} + c(\mathbf{y}, \phi) \right] \\ &= \frac{1}{\phi} \frac{\partial \mathbf{z}}{\partial \boldsymbol{\theta}} \left(\mathbf{y} - \frac{\partial \eta}{\partial \mathbf{z}} \right) \end{aligned} \quad (50)$$

we have

$$F(\boldsymbol{\theta}) = \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x}), \mathbf{y} \sim P(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta})} \left[\frac{\partial \mathbf{z}}{\partial \boldsymbol{\theta}} \frac{(\mathbf{y} - \eta'(\mathbf{z}))}{\phi} \frac{(\mathbf{y} - \eta'(\mathbf{z}))^T}{\phi} \frac{\partial \mathbf{z}^T}{\partial \boldsymbol{\theta}} \right] \quad (51)$$

$$= \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[\frac{\partial \mathbf{z}}{\partial \boldsymbol{\theta}} \mathbb{E}_{\mathbf{y} \sim P(\mathbf{y}|\mathbf{x}; \boldsymbol{\theta})} \left[\frac{(\mathbf{y} - \eta'(\mathbf{z}))}{\phi} \frac{(\mathbf{y} - \eta'(\mathbf{z}))^T}{\phi} \middle| \mathbf{x}; \boldsymbol{\theta} \right] \frac{\partial \mathbf{z}^T}{\partial \boldsymbol{\theta}} \right] \quad (52)$$

$$= \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[\frac{\partial \mathbf{z}}{\partial \boldsymbol{\theta}} \frac{\text{Cov}[\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}]}{\phi^2} \frac{\partial \mathbf{z}^T}{\partial \boldsymbol{\theta}} \right] \quad (53)$$

$$= \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[\frac{\partial \mathbf{z}}{\partial \boldsymbol{\theta}} \frac{\text{diag}(f'(\mathbf{z}))}{\phi} \frac{\partial \mathbf{z}^T}{\partial \boldsymbol{\theta}} \right] \quad (54)$$

where we recognise the covariance matrix is diagonal:

$$\text{Cov}[\mathbf{y}|\mathbf{x}; \boldsymbol{\theta}] = \text{diag}(\text{Var}(y_1|\mathbf{x}; \boldsymbol{\theta}), \dots, \text{Var}(y_{n_e}|\mathbf{x}; \boldsymbol{\theta}))$$

To analyse the approximate KL divergence resulting from the simple case of perturbing individual parameters of a single-layer DANN, we only need to consider the diagonal entries of the Fisher.

$$D_{\text{KL}}[P_{\boldsymbol{\theta}} \| P_{\boldsymbol{\theta} + \delta_{\bar{\theta}}}] \approx \frac{1}{2} \delta_{\bar{\theta}}^T \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[\frac{\partial \mathbf{z}}{\partial \boldsymbol{\theta}} \frac{\text{diag}(f'(\mathbf{z}))}{\phi} \frac{\partial \mathbf{z}^T}{\partial \boldsymbol{\theta}} \right] \delta_{\bar{\theta}} \quad (55)$$

$$= \frac{\delta^2}{2\phi} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[\frac{\partial \mathbf{z}}{\partial \boldsymbol{\theta}} \text{diag}(f'(\mathbf{z})) \frac{\partial \mathbf{z}^T}{\partial \boldsymbol{\theta}} \right] \quad (56)$$

$$= \frac{\delta^2}{2\phi} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[(f'(\mathbf{z}))^T \odot \frac{\partial \mathbf{z}}{\partial \boldsymbol{\theta}} \frac{\partial \mathbf{z}^T}{\partial \boldsymbol{\theta}} \right] \quad (57)$$

$$= \frac{\delta^2}{2\phi} \sum_k^{n_e} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_k) \left(\frac{\partial z_k}{\partial \bar{\theta}} \right)^2 \right] \quad (58)$$

where $\delta_{\tilde{\theta}}$ represents a 1-hot vector corresponding to $\tilde{\theta}$, multiplied by a scalar δ .

D.2 DERIVATIVES

Here we provide derivatives for DANN layer activations with respect to the different parameter groups. The equations for the layer activation can be written

$$\mathbf{z} = \frac{\mathbf{g}}{\gamma} \odot (\mathbf{z}^E - \mathbf{W}^{EI} \mathbf{z}^I) + \beta$$

where $\mathbf{z}^E = \mathbf{W}^{EE} \mathbf{x}$ $\mathbf{z}^I = \mathbf{W}^{IE} \mathbf{x}$ $\gamma = \mathbf{W}^{EI} (e^{\alpha} \odot \mathbf{z}^I)$ (59)

Note $\frac{\partial z_k}{\partial w_{ij}^{EE}} = \frac{\partial z_k}{\partial w_{ij}^{EI}} = 0$ for $k \neq i$.

$$\frac{\partial z_i}{\partial w_{ij}^{EE}} = \frac{\partial}{\partial w_{ij}^{EE}} \left(\frac{g_i}{\gamma_i} (z_i^E - (\mathbf{W}^{EI} \mathbf{z}^I)_i) + \beta_i \right)$$
 (60)

$$= \frac{g_i}{\gamma_i} \frac{\partial}{\partial w_{ij}^{EE}} (z_i^E)$$
 (61)

$$= \frac{g_i}{\gamma_i} x_j$$
 (62)

$$\frac{\partial z_i}{\partial w_{ij}^{EI}} = \frac{\partial}{\partial w_{ij}^{EI}} \left(\frac{g_i}{\gamma_i} (z_i^E - (\mathbf{W}^{EI} \mathbf{z}^I)_i) + \beta_i \right)$$
 (63)

$$= -\frac{g_i}{\gamma_i^2} \frac{\partial \gamma_i}{\partial w_{ij}^{EI}} (z_i^E - (\mathbf{W}^{EI} \mathbf{z}^I)_i) - \frac{g_i}{\gamma_i} z_j^I$$
 (64)

$$= -\frac{g_i}{\gamma_i^2} e^{\alpha_j} z_j^I (z_i^E - (\mathbf{W}^{EI} \mathbf{z}^I)_i) - \frac{g_i}{\gamma_i} z_j^I$$
 (65)

$$= -\frac{g_i}{\gamma_i} z_j^I \left(\frac{e^{\alpha_j}}{\gamma_i} (z_i^E - (\mathbf{W}^{EI} \mathbf{z}^I)_i) + 1 \right)$$
 (66)

$$= -\sum_k \frac{g_i}{\gamma_i} w_{jk}^{IE} x_k \left(\frac{e^{\alpha_j}}{\gamma_i} (z_i^E - (\mathbf{W}^{EI} \mathbf{z}^I)_i) + 1 \right)$$
 (67)

In contrast $\frac{\partial z_k}{\partial w_{ij}^{EI}}, \frac{\partial z_k}{\partial \alpha_j} \neq 0$ for $k \neq i$.

$$\frac{\partial z_k}{\partial w_{ij}^{IE}} = \frac{\partial}{\partial w_{ij}^{IE}} \left(\frac{g_k}{\gamma_k} (z_k^E - (\mathbf{W}^{EI} \mathbf{z}^I)_k) + \beta_k \right)$$
 (68)

$$= -\frac{g_k}{\gamma_k^2} \frac{\partial \gamma_k}{\partial w_{ij}^{IE}} (z_k^E - (\mathbf{W}^{EI} \mathbf{z}^I)_k) - \frac{g_k}{\gamma_k} w_{ki}^{EI} x_j$$
 (69)

$$= -\frac{g_k}{\gamma_k^2} e^{\alpha_i} w_{ki}^{EI} x_j (z_k^E - (\mathbf{W}^{EI} \mathbf{z}^I)_k) - \frac{g_k}{\gamma_k} w_{ki}^{EI} x_j$$
 (70)

$$= -\frac{g_k}{\gamma_k} w_{ki}^{EI} x_j \left(\frac{e^{\alpha_i}}{\gamma_k} (z_k^E - (\mathbf{W}^{EI} \mathbf{z}^I)_k) + 1 \right)$$
 (71)

$$\frac{\partial z_i}{\partial \alpha_j} = \frac{\partial}{\partial \alpha_j} \left(\frac{g_i}{\gamma_i} (z_i^E - (\mathbf{W}^{\text{EI}} \mathbf{z}^{\text{I}})_i) + \beta_i \right) \quad (72)$$

$$= -\frac{g_i}{\gamma_i^2} \frac{\partial \gamma_i}{\partial \alpha_j} (z_i^E - (\mathbf{W}^{\text{EI}} \mathbf{z}^{\text{I}})_i) \quad (73)$$

$$= -\frac{g_i}{\gamma_i^2} w_{ij}^{\text{EI}} e^{\alpha_j} z_j^{\text{I}} (z_i^E - (\mathbf{W}^{\text{EI}} \mathbf{z}^{\text{I}})_i) \quad (74)$$

$$= -\sum_k^d \frac{g_i}{\gamma_i} w_{ij}^{\text{EI}} w_{jk}^{\text{IE}} x_k \frac{e^{\alpha_j}}{\gamma_i} (z_i^E - (\mathbf{W}^{\text{EI}} \mathbf{z}^{\text{I}})_i) \quad (75)$$

$$\frac{\partial z_i}{\partial g_i} = \frac{1}{\gamma_i} (z_i^E - (\mathbf{W}^{\text{EI}} \mathbf{z}^{\text{I}})_i) \quad (76)$$

$$\frac{\partial z_i}{\partial b_i} = 1 \quad (77)$$

$$(78)$$

D.3 APPROXIMATE KL DIVERGENCE FOR WEIGHT UPDATES

If we consider an update to an element ij of \mathbf{W}^{EE} the approximate KL divergence is

$$\begin{aligned} D_{\text{KL}}[P_{\boldsymbol{\theta}} \parallel P_{\boldsymbol{\theta} + \delta \mathbf{W}_{ij}^{\text{EE}}}] &\approx \frac{\delta^2}{2\phi} \sum_k^{n_e} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_k) \left(\frac{\partial z_k}{\partial w_{ij}^{\text{EE}}} \right)^2 \right] \\ &= \frac{\delta^2}{2\phi} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_i) \left(\frac{g_i}{\gamma_i} x_j \right)^2 \right] \end{aligned} \quad (79)$$

as $\frac{\partial z_k}{\partial w_{ij}^{\text{EE}}} = 0$ for $k \neq i$.

In contrast, for an update to an element ij of \mathbf{W}^{IE} we sum over n_e terms, as $\frac{\partial z_k}{\partial w_{ij}^{\text{IE}}} \neq 0$ for $k \neq i$.

$$\begin{aligned} D_{\text{KL}}[P_{\boldsymbol{\theta}} \parallel P_{\boldsymbol{\theta} + \delta \mathbf{W}_{ij}^{\text{IE}}}] &\approx \frac{\delta^2}{2\phi} \sum_k^{n_e} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_k) \left(\frac{\partial z_k}{\partial w_{ij}^{\text{IE}}} \right)^2 \right] \\ &= \frac{\delta^2}{2\phi} \sum_k^{n_e} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_k) \left(-\frac{g_k}{\gamma_k} w_{ki}^{\text{EI}} x_j a_{ki} \right)^2 \right] \\ &= \frac{\delta^2}{2\phi} \sum_k^{n_e} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_k) \left(\frac{g_k}{\gamma_k} x_j \right)^2 (w_{ki}^{\text{EI}} a_{ki})^2 \right] \end{aligned} \quad (80)$$

where $a_{kj} = \frac{e^{\alpha_j}}{\gamma_k} (z_k^E - (\mathbf{W}^{\text{EI}} \mathbf{z}^{\text{I}})_k) + 1$.

For an update $\delta \mathbf{w}_{ij}^{\text{EI}}$, while $\frac{\partial z_k}{\partial w_{ij}^{\text{EI}}} = 0$ for $k \neq i$, the derivative contains a z_j^{I} term, so there is instead a squared sum over d terms.

$$\begin{aligned}
 \text{D}_{\text{KL}}[P_{\boldsymbol{\theta}} \parallel P_{\boldsymbol{\theta} + \delta \mathbf{w}_{ij}^{\text{EI}}}] &\approx \frac{\delta^2}{2\phi} \sum_k^{n_e} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_k) \left(\frac{\partial z_k}{\partial w_{ij}^{\text{EI}}} \right)^2 \right] \\
 &= \frac{\delta^2}{2\phi} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_i) \left(\frac{\partial z_i}{\partial w_{ij}^{\text{EI}}} \right)^2 \right] \\
 &= \frac{\delta^2}{2\phi} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_i) \left(-\frac{g_i}{\gamma_i} z_j^{\text{I}} a_{ij} \right)^2 \right] \\
 &= \frac{\delta^2}{2\phi} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_i) (z_j^{\text{I}})^2 \left(\frac{g_i}{\gamma_i} \right)^2 (a_{ij})^2 \right] \\
 &= \frac{\delta^2}{2\phi} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_i) \left(\sum_n^d w_{jn}^{\text{IE}} x_n \right)^2 \left(\frac{g_i}{\gamma_i} \right)^2 (a_{ij})^2 \right] \\
 &= \frac{\delta^2}{2\phi} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_i) \left(\sum_n^d (w_{jn}^{\text{IE}})^2 (x_n)^2 + \sum_{n \neq m}^d w_{jn}^{\text{IE}} w_{jm}^{\text{IE}} x_n x_m \right) \left(\frac{g_i}{\gamma_i} \right)^2 (a_{ij})^2 \right] \\
 &= \frac{\delta^2}{2\phi} \sum_n^d \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_i) (w_{jn}^{\text{IE}})^2 (x_n)^2 \left(\frac{g_i}{\gamma_i} \right)^2 (a_{ij})^2 \right] \\
 &\quad + \frac{\delta^2}{2\phi} \sum_{n \neq m}^d \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_i) w_{jn}^{\text{IE}} w_{jm}^{\text{IE}} x_n x_m \left(\frac{g_i}{\gamma_i} \right)^2 (a_{ij})^2 \right]
 \end{aligned} \tag{81}$$

Finally, for alpha

$$\begin{aligned}
 \text{D}_{\text{KL}}[P_{\boldsymbol{\theta}} \parallel P_{\boldsymbol{\theta} + \delta \alpha_i}] &\approx \frac{\delta^2}{2\phi} \sum_k^{n_e} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_k) \left(\frac{\partial z_k}{\partial \alpha_i} \right)^2 \right] \\
 &= \frac{\delta^2}{2\phi} \sum_k^{n_e} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_k) \left(-\sum_j^d \frac{g_k}{\gamma_k} w_{ki}^{\text{EI}} w_{ij}^{\text{IE}} x_j \frac{e^{\alpha_i}}{\gamma_k} (z_k^{\text{E}} - (\mathbf{W}^{\text{EI}} \mathbf{z}^{\text{I}})_k)^2 \right)^2 \right] \\
 &= \frac{\delta^2}{2\phi} \sum_k^{n_e} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_k) \left(-\sum_j^d \frac{g_k}{\gamma_k} w_{ki}^{\text{EI}} w_{ij}^{\text{IE}} x_j (a_{ki} - 1) \right)^2 \right] \\
 &= \frac{\delta^2}{2\phi} \sum_k^{n_e} \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_k) \left(\sum_j^d \left(\frac{g_k}{\gamma_k} w_{ki}^{\text{EI}} w_{ij}^{\text{IE}} x_j (a_{ki} - 1) \right)^2 \right. \right. \\
 &\quad \left. \left. + \sum_{n \neq m}^d \left(\left(\frac{g_k}{\gamma_k} w_{ki}^{\text{EI}} \right)^2 w_{in}^{\text{IE}} x_n w_{im}^{\text{IE}} x_m (a_{ki} - 1)^2 \right) \right) \right] \\
 &= \frac{\delta^2}{2\phi} \sum_k^{n_e} \sum_j^d \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_k) \left(\frac{g_k}{\gamma_k} w_{ki}^{\text{EI}} w_{ij}^{\text{IE}} x_j (a_{ki} - 1) \right)^2 \right] \\
 &\quad + \frac{\delta^2}{2\phi} \sum_k^{n_e} \sum_{n \neq m}^d \mathbb{E}_{\mathbf{x} \sim P(\mathbf{x})} \left[f'(z_k) \left(\frac{g_k}{\gamma_k} w_{ki}^{\text{EI}} \right)^2 w_{in}^{\text{IE}} x_n w_{im}^{\text{IE}} x_m (a_{ki} - 1)^2 \right]
 \end{aligned} \tag{82}$$

E ALGORITHMS

Here we provide pseudo-code for implementation details. Please see the following link for code: [Learning_to_live_with_Dales_principle](#)

E.1 PARAMETER INITIALIZATION

Algorithm 1 Parameter initialization for DANNs

```

for layer  $L$  do
  require  $n_e, n_i, d$ 
   $\mathbf{W}^{\text{EE}} \sim \exp(\lambda^{\text{E}})$ 
  if  $n_i = 1$ 
     $\mathbf{W}^{\text{IE}} \leftarrow \frac{1}{n_e} \sum_{j=1}^{n_e} \mathbf{w}_j^{\text{EE}}$ 
     $\mathbf{W}^{\text{EI}} \leftarrow \mathbb{I}$ 
  else:
     $\mathbf{W}^{\text{IE}} \sim \exp(\lambda^{\text{E}})$ 
     $\mathbf{W}^{\text{EI}} \leftarrow \mathbb{I}/n_i$ 
  end if
   $\alpha \leftarrow \mathbb{I} \cdot \log(\frac{\sqrt{2\pi-1}}{\sqrt{d}})$ 
   $\mathbf{g}, \beta \leftarrow \mathbb{I}$ 
end for

```

Where number of excitatory output units is n_e , number of inhibitory units n_i , and input dimensionality d and $\lambda^{\text{E}} = \sqrt{d(2\pi-1)}/\sqrt{2\pi}$.

E.2 PARAMETER UPDATES

For DANN parameter updates we used the algorithms detailed below. Note that gradients were corrected as detailed in Section 4 and see Algorithm 3.

All below algorithms are computed using the loss gradients $\nabla\theta$ of parameter θ , in a given model computed on a minibatch sample.

Algorithm 2 Parameter updates

```

Require learning rate  $\eta$ , updates  $\Delta\theta$ 
for each layer  $l$  do
   $\mathbf{W}^{\text{EE}} \leftarrow \mathbf{W}^{\text{EE}} - \eta\Delta\mathbf{W}^{\text{EE}}$ 
   $\mathbf{W}^{\text{IE}} \leftarrow \mathbf{W}^{\text{IE}} - \eta\Delta\mathbf{W}^{\text{IE}}$ 
   $\mathbf{W}^{\text{EI}} \leftarrow \mathbf{W}^{\text{EI}} - \eta\Delta\mathbf{W}^{\text{EI}}$ 
   $\alpha \leftarrow \alpha - \eta\Delta\alpha$ 
   $\mathbf{g} \leftarrow \mathbf{g} - \eta\Delta\mathbf{g}$ 
   $\beta \leftarrow \beta - \eta\Delta\beta$ 
   $\mathbf{W}^{\text{EE}} \leftarrow \max(\mathbf{W}^{\text{EE}}, 0)$ 
   $\mathbf{W}^{\text{IE}} \leftarrow \max(\mathbf{W}^{\text{IE}}, 0)$ 
   $\mathbf{W}^{\text{EI}} \leftarrow \max(\mathbf{W}^{\text{EI}}, 0)$ 
   $\mathbf{g} \leftarrow \max(\mathbf{g}, 0)$ 
end for

```

E.3 DANN GRADIENT CORRECTION ALGORITHMS

For the majority of experiments we scaled gradients using the heuristic correction terms derived in Section 4 (and see Appendix D). In this case we applied the following algorithm before Algorithm 2.

Algorithm 3 DANN gradient correction

```

for each layer  $l$  do
  require  $n_e, n_i, d$ 
   $\Delta \mathbf{W}^{EE} \leftarrow \nabla \mathbf{W}^{EE}$ 
   $\Delta \mathbf{W}^{IE} \leftarrow \frac{1}{\sqrt{n_e}} \nabla \mathbf{W}^{IE}$ 
   $\Delta \mathbf{W}^{EI} \leftarrow \frac{1}{d} \nabla \mathbf{W}^{EI}$ 
   $\Delta \alpha \leftarrow \frac{1}{d\sqrt{n_e}} \nabla \alpha$ 
   $\Delta \mathbf{g} \leftarrow \nabla \mathbf{g}$ 
   $\Delta \beta \leftarrow \nabla \beta$ 
end for

```

We also tested that our heuristic correction factors approximated gradient multiplication by the diagonal of F_t^{-1} for each layer (see Figure 2).

Algorithm 4 Gradient correction by approximation of $\text{diag}(\mathbf{F}^{-1})$

```

Require learning rate  $\eta$ , fisher momentum  $k$ , fisher learning rate  $\lambda$ 
for each batch  $(x_t, y_t)$  do
  Compute  $p = \text{softmax}(z)$ 
  Compute cross entropy loss  $L(p, y_t)$ 
  for each layer  $l$  do
     $\nabla \theta^l \leftarrow \frac{\partial L}{\partial \theta^l}$ 
  end for
  Sample  $\hat{y} \sim \text{Categorical}(p)$ 
  Compute cross entropy loss  $L(p, \hat{y})$ 
  for each layer  $l$  do
     $\hat{F} = \frac{1}{\lambda \cdot |\text{batch}|} \sum_i^{|\text{batch}|} \left( \frac{\partial L}{\partial \theta^l} \right)_i^2$ 
     $F_t = k\hat{F} + (1 - k)F_{t-1}$ 
     $F_t^{-1} = 1/F_t$ 
     $F_t^* = F_t^{-1} \cdot 1/\|\nabla \theta^l\|$  where  $F^{\mathbf{W}^{EE}}$  is the elements of  $F$  corresponding to  $\mathbf{W}^{EE}$ 
     $\Delta \theta^l \leftarrow F_t^* \nabla \theta^l$ 
  end for
end for

```

Here we note this update can be considered very rough diagonal approximation to natural gradient descent. In addition, various efficient approximations to natural gradient descent that have been utilized such as KFAC Martens & Grosse (2015) could not be applied due to the structure of DANNs, as the mathematical assumptions of KFAC, which were made for feedforward networks with activations as matrix multiplications, do not apply.

E.4 LEARNING RATE SCALING AND GRADIENT NORMALISATION

We also tested whether constraining the gradient norm and scaling the learning rate based on parameter clipping improved DANN performance. For these experiments we applied the following algorithms.

Algorithm 5 Gradient normalisation

```

for each layer  $\ell$  do
  Require  $\nabla \theta^\ell, M$ 
  if  $\|\nabla \theta^\ell\|_2 > M$ :
     $\Delta \theta^\ell \leftarrow M \cdot \frac{\nabla \theta^\ell}{\|\nabla \theta^\ell\|_2}$ 
  else:
     $\Delta \theta^\ell \leftarrow \nabla \theta^\ell$ 
  end if
end for

```

Algorithm 6 Learning rate scaling

```

for each layer  $\ell$  do
  Require  $\nabla\theta^\ell, M, \xi$ 
   $i \leftarrow 1$ 
   $c \leftarrow 0$ 
  while  $c < M$ :
     $\eta \leftarrow \xi_i$ 
     $c \leftarrow \text{CosineSimilarity}(\max(0, \theta^\ell - \eta\nabla\theta^\ell), \theta^\ell - \eta\nabla\theta^\ell)$ 
     $i \leftarrow i + 1$ 
  end for

```

The learning rate scaling method temporarily reduces the learning rate whenever parameter clipping would reduce the cosine of the angle, made between the gradient and actual updates, below a certain constraint. For any optimization problems caused by actual clipped updates not following the gradient, learning rate scaling is a principled way of following the direction of the gradient.

We also note, this technique can be generally applied to any other model which is constrained so that it cannot have updates freely follow gradient descent. If the constrained parameter space is an open subset of euclidean space, and we allow the learning rate to be arbitrarily small (Algorithm 6 with $\lim_{i \rightarrow \infty} \xi_i = 0$), updates will always follow the direction of the gradient.