



# BEAN: Interpretable and Efficient Learning With Biologically-Enhanced Artificial Neuronal Assembly Regularization

Yuyang Gao<sup>1</sup>, Giorgio A. Ascoli<sup>2</sup> and Liang Zhao<sup>1\*</sup>

<sup>1</sup> Department of Information Sciences and Technology, George Mason University, Fairfax, VA, United States, <sup>2</sup> Department of Bioengineering, Center for Neural Informatics, Krasnow Institute for Advanced Study, George Mason University, Fairfax, VA, United States

## OPEN ACCESS

### Edited by:

Jeffrey L. Krichmar,  
University of California, Irvine,  
United States

### Reviewed by:

Marc De Kamps,  
University of Leeds, United Kingdom  
Sohail Kolouri,  
HRL Laboratories, United States

### \*Correspondence:

Liang Zhao  
lzha09@gmu.edu

**Received:** 29 May 2020

**Accepted:** 05 May 2021

**Published:** 01 June 2021

### Citation:

Gao Y, Ascoli GA and Zhao L (2021)  
BEAN: Interpretable and Efficient  
Learning With Biologically-Enhanced  
Artificial Neuronal Assembly  
Regularization.  
*Front. Neurobot.* 15:567482.  
doi: 10.3389/fnbot.2021.567482

Deep neural networks (DNNs) are known for extracting useful information from large amounts of data. However, the representations learned in DNNs are typically hard to interpret, especially in dense layers. One crucial issue of the classical DNN model such as multilayer perceptron (MLP) is that neurons in the same layer of DNNs are conditionally independent of each other, which makes co-training and emergence of higher modularity difficult. In contrast to DNNs, biological neurons in mammalian brains display substantial dependency patterns. Specifically, biological neural networks encode representations by so-called neuronal assemblies: groups of neurons interconnected by strong synaptic interactions and sharing joint semantic content. The resulting population coding is essential for human cognitive and mnemonic processes. Here, we propose a novel Biologically Enhanced Artificial Neuronal assembly (BEAN) regularization<sup>1</sup> to model neuronal correlations and dependencies, inspired by cell assembly theory from neuroscience. Experimental results show that BEAN enables the formation of interpretable neuronal functional clusters and consequently promotes a sparse, memory/computation-efficient network without loss of model performance. Moreover, our few-shot learning experiments demonstrate that BEAN could also enhance the generalizability of the model when training samples are extremely limited.

**Keywords:** deep learning-artificial neural network (DL-ANN), neuronal assemblies, neuronal correlations, regularization, explainability, interpretability, explainable AI, representation learning

## 1. INTRODUCTION

Deep neural networks (DNNs) are known for extracting useful information from a large amount of data (Bengio et al., 2013). Despite the success and popularity of DNNs in a wide variety of fields, including computer vision (Krizhevsky et al., 2012; He et al., 2016) and natural language processing (Collobert and Weston, 2008; Young et al., 2018), there are still many drawbacks and limitations of modern DNNs, including lack of interpretability (Zhang and Zhu, 2018), the requirement of large data (Kimura et al., 2018), and post selection on complex model architecture (Zheng and Weng, 2016a,b). Specifically, the representations learned in DNNs are typically hard to interpret, especially in dense (fully connected) layers. Despite recent attempts to build intrinsically

<sup>1</sup>The code is available at <https://github.com/YuyangGao/BEAN>.

more interpretable convolutional units (Sabour et al., 2017; Zhang and Zhu, 2018), the exploration of learned representations in the dense layer has remained limited. In fact, dense layers are the fundamental and critical component of most state-of-the-art DNNs, which are typically used for the late stage of the network's computation, akin to the inference and decision-making processes (Krizhevsky et al., 2012; Simonyan and Zisserman, 2014; He et al., 2016). Thus improving the interpretability of the dense layer representation is crucial if we are to fully understand and exploit the power of DNNs.

However, interpreting the representations learned in dense layers of DNNs is typically a very challenging task. One crucial issue of the classical DNN model such as multilayer perceptron (MLP) is that neurons in the same layer of DNNs are conditionally independent of each other, as dense layers in MLP are typically activated by all-to-all feed-forward neuron activity and trained by all-to-all feedback weight adjustment. In this comprehensively “vertical” connectivity, every node is independent and abstracted “out of the context” of the other nodes. This issue limits the analysis of the representation learned in DNNs to single-unit level, as opposed to the higher modularity in principle afforded by neuron population coding. Moreover, recent studies on single unit importance seem to suggest that individually selective units may have little correlation with overall network performance (Morcos et al., 2018; Zhou et al., 2018). Specifically, Morcos et al. (2018) and Zhou et al. (2018) conducted unit-level ablation experiments on CNNs trained on large scale image datasets and found that ablating any individual unit does not hurt overall classification accuracy.

On the other hand, understanding the complex patterns of neuron correlations in biological neural networks (BNNs) has long been a subject of intense interest for neuroscience researchers. Circuitry blueprints in the real brain are “filtered” by the physical requirements of axonal projections and the consequent need to minimize cable while maximizing connections (Ropireddy et al., 2011). One could naively expect that the non-all-to-all limitations imposed in natural neural systems would be detrimental to their computational power. Instead, it makes them superiorly efficient and allows cell assemblies to emerge. Neuronal assemblies or cell assemblies (Hebb, 1949) can be described as groups of neurons interconnected by strong synaptic interactions and sharing joint semantic content. The resulting population coding is essential for human cognitive and mnemonic processes (Braitenberg, 1978).

In this paper, we bridge such a crucial gap between DNNs and BNNs by modeling the neuron correlations within each layer of DNNs. Leveraging biologically inspired learning rules in neuroscience and graph theory, we propose a novel Biologically-Enhanced Artificial Neuronal assembly (BEAN) regularization that can enforce dependencies among neurons in dense layers of DNNs without substantially altering the conventional architecture. The resultant advantages are threefold:

- **Enhancing interpretability and modularity at the neuron population level.** Modeling neural correlations and dependencies allows us to better interpret and visualize

the learned representation in hidden layers at the neuron population level instead of the single neuron level. Both qualitative and quantitative analyses show that BEAN enables the formations of identifiable neuronal assembly patterns in the hidden layers, enhancing the modularity and interpretability of the DNN representations.

- **Promoting jointly sparse and efficient encoding of rich semantic correlation among neurons.** Here, we show that BEAN can promote jointly sparse and efficient encoding of rich semantic correlation among neurons in DNNs similar to connection patterns in BNNs. BEAN enables the model to parsimoniously leverage available neurons and possible connections through modeling structural correlation, yielding both connection-level and neuron-level sparsity in the dense layers. Experimental results show that BEAN not only enables the formation of neuronal functional clusters that encode rich semantic correlation, but also allows the model to achieve state-of-the-art memory/computational efficiency without loss of model performance.
- **Improving model generalizability with few training samples.** Humans and animals can learn and generalize to new concepts with just a few trials of learning, while DNNs generally perform poorly on such tasks. Current few-shot learning techniques in deep learning still rely heavily on a large amount of additional knowledge to work well. For example, transfer-learning-based methods typically leverage a model pre-trained with a large amount of data (Socher et al., 2013; Xian et al., 2018), and meta-learning-based methods require a large number of additional side tasks (Finn et al., 2017; Snell et al., 2017). Here we explore BEAN with a substantially more challenging *few-shot learning from scratch* task first studied by Kimura et al. (2018), where no additional knowledge is provided aside from a few training observations. Extensive experiments show that BEAN has a significant advantage in improving model generalizability over conventional techniques.

## 2. BIOLOGICALLY-ENHANCED ARTIFICIAL NEURONAL ASSEMBLY REGULARIZATION

This section describes the overall objective of Biologically-Enhanced Artificial Neuronal Assembly (BEAN) regularization as well as the implementation of BEAN on DNNs, as *Layer-wise Neuron Correlation and Co-activation Divergence* to model the implicit dependencies between neurons within the same layer.

### 2.1. Layer-Wise Neuron Co-activation Divergence

Due to the physical restrictions imposed by dendrites and axons (Rivera-Alba et al., 2014) and for energy efficiency, biological neural systems are “parsimonious” and can only afford to form a limited number of connections between neurons. The neuron connectivity patterns of BNNs are intertwined with their activation patterns based on the principle of “*Cells that fire together wire together*,” which is known as **cell assembly theory**. It explains and relates to several characteristics and advantages

of BNN architecture such as modularity (Peyrache et al., 2010), efficiency, and generalizability, that are just the aspects in which the current DNNs are usually struggling (LeCun et al., 2015). To take advantage of the beneficial architectural features in BNNs and overcome the existing drawbacks of DNNs, we propose the Biologically-Enhanced Artificial Neuronal assembly (BEAN) regularization. BEAN ensures neurons which “wire” together with a high outgoing weight correlation also “fire” together with small divergence in terms of their activation patterns.

An example of the artificial neuronal assembly achieved by our method can be seen in **Figure 1D**. The regularization is formulated as follows:

$$L_c^{(l)} = 1/(SN_l^2) \sum_s \sum_i \sum_j A_{ij}^{(l)} \times d(H_{s,i}^{(l)}, H_{s,j}^{(l)}) \quad (1)$$

where  $L_c$  is the regularization loss; the term  $A_{ij}^{(l)}$  characterizes the wiring strength (the higher value, the stronger connection) between two neurons  $i$  and  $j$  within layer  $l$ ; the term  $d(H_{s,i}^{(l)}, H_{s,j}^{(l)})$  models the divergence of firing patterns (the higher value, the more different the firing) between two neurons  $i$  and  $j$  on input sample  $s$ . Thus, by multiplying these two functions, we penalize those neurons with strong connectivity but high activation divergence, in line with the principles of cell assembly theory.  $S$  is the total number of input samples while  $N_l$  is the total number of hidden neurons in layer  $l$ .

Specifically,  $A_{ij}^{(l)}$  defines the connectivity relation among neuron  $i$  and neuron  $j$  in DNN, which is instantiated by our newly proposed “Layer-wise Neuron Correlation” and will be elaborated in sections 2.2 and 2.3. On the other hand, to model the “co-firing” correlation,  $d(H_{s,i}^{(l)}, H_{s,j}^{(l)})$  is defined as “Layer-wise Neuron Co-activation Divergence” which denotes the difference in the activation patterns in  $l$ th layer between  $H_{s,i}^{(l)}$  and  $H_{s,j}^{(l)}$  of neuron  $i$  and neuron  $j$ , respectively. Here  $H_{s,i}^{(l)}$  represents the activation of neuron  $i$  in layer  $l$  for a given input sample  $s$ . The function  $d(x, y)$  can be a common divergence metric such as absolute difference or square difference. In this study, we show the results for a square difference in the Experimental Study Section; the absolute difference results follow a similar trend.

**Model Training:** The general objective function of training a DNN model along with the proposed regularization on fully connected layer  $l$  can be written as:  $L = L_{DNN} + \alpha L_c^{(l)}$ , where  $L_{DNN}$  represents the general deep learning model training loss and the hyper-parameter  $\alpha$  controls the relative strength of the regularization.

Equation (1) can be optimized with backpropagation (Rumelhart et al., 1988) using the chain rule:

$$\frac{\partial L_c^{(l)}}{\partial W^{(l+1)}} = \frac{\partial A^{(l)}}{\partial W^{(l+1)}} D^{(l)}, \quad \frac{\partial L_c^{(l)}}{\partial W^{(l)}} = A^{(l)} \frac{\partial D^{(l)}}{\partial H^{(l)}} \frac{\partial H^{(l)}}{\partial W^{(l)}}, \dots \quad (2)$$

where  $D^{(l)} \in \mathbb{R}^{S \times N_l \times N_l}$  of which each element is  $D_{s,i,j}^{(l)} = d(H_{s,i}^{(l)}, H_{s,j}^{(l)})$ .

**Remark 1.** BEAN regularization has several strengths. First, it enforces interpretable neuronal assemblies without the need to introduce sophisticated handcrafted designs into the architecture,

which is justified later in section 3.1. In addition, modeling the neuron correlations and dependencies further results in sparse and efficient connectivity in dense layers, which substantially reduced the computation/memory cost of the model, as shown in section 3.2. Besides, the encoding of rich semantic correlation among neurons may improve the generalizability of the model when insufficient data and knowledge are provided, which is demonstrated later in section 3.3. Finally, the Layer-wise Neuron Correlation can be efficiently computed with matrix operations, as per Equations (5, 7), which enables modern GPUs to boost up the speed during model training. In practice, we observe negligible run time overhead of the addition computation needed for BEAN regularization.

## 2.2. The First-Order Layer-Wise Neuron Correlation

This section introduces the formulation of the layer-wise neuron correlation  $A_{ij}^{(l)}$  between any pair of neurons  $i$  and  $j$ .

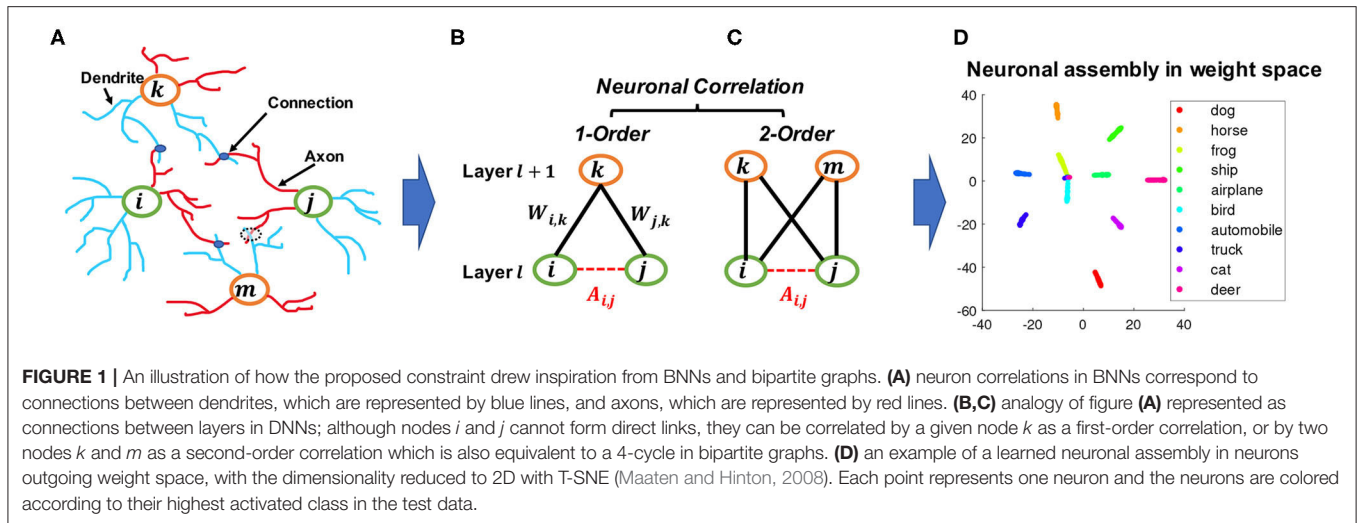
In the human brain, the correlation between two neurons depends on the wiring between them (Buzsáki, 2010) and hence is typically treated as a binary value in BNN studies, with “1” indicating the presence of a connection and “0” the absence, so the correlation among a group of neurons can be represented by the corresponding adjacency matrix. Although there is typically no direct connection between neurons within the same layer of DNNs, it is possible to model neuron correlations based on their connectivity patterns to the next layer. This resembles a common approach in network science, where it is useful to consider the relationships between nodes based on their common neighbors in addition to their direct connections. One classic concept widely used to describe such a pattern is called *triadic closure* (Granovetter, 1973). As shown in **Figure 1B**, triadic closure can be interpreted here as a property among three nodes  $i$ ,  $j$ , and  $k$ , such that if connections exist between  $i - k$  and  $j - k$ , there is also a connection between  $i - j$ .

We take this scheme a step further to model the correlations between neurons within the same layer by their connections to the neurons in the next layer. This can be considered loosely analogous to the degree of similarity of the axonal connection pattern of biological neurons in BNNs (Rees et al., 2017). To simulate the relative strength of such connections in DNNs, we introduce a function  $f(\cdot)$  that converts the actual weights into a relative connectivity strength. Suppose matrix  $W^{(l+1)} \in \mathbb{R}^{N_l \times N_{l+1}}$  represents all the weights between neurons in layers  $l$  and  $l + 1$  in DNNs, where  $N_l$  and  $N_{l+1}$  represent the numbers of neurons, respectively. The relative connectivity strength can be estimated by the following equation<sup>2</sup>:

$$f(W^{(l+1)}) = |\tanh(\gamma W^{(l+1)})| \quad (3)$$

where  $|\cdot|$  represents the element-wise absolute operator;  $\tanh(\cdot)$  represents the element-wise hyperbolic tangent function; and  $\gamma$  is a scalar that controls the curvature of the hyperbolic tangent function. The values of  $f(W^{(l+1)}) \in \mathbb{R}^{N_l \times N_{l+1}}$  will all be positive and in the range of  $[0, 1]$  with the value simulating

<sup>2</sup>Similar to the ReLU activation function, our formulation introduces a non-differentiable point at zero; we follow the conventional setting by using the sub-gradient for model optimization.



the relative connectivity strength of the corresponding synapse between neurons.

Although there can be positive and negative weights in DNNs, our assumption on connection strength follows the typical way of BNN studies, which measures the presence and absence of the connection as mentioned above. Moreover, since DNNs require continuous values instead of discrete values to make the function differentiable for optimization, we further use Equation (3) to convert the concept of the presence/absence of the connections to the relative strength of the connections. More specifically, the difference is that instead of treating connection to be either “1” (indicating the presence of a connection) or “0” (indicating the absence of the connection), we treat the output of Equation (3) as the strength of that connection, where high values (i.e., close to “1”) indicate the presence of a strong connection and low values (i.e., close to “0”) indicate weak or no connection.

Based on this, we can now give the definition for the *layer-wise first-order neuron correlation* as:

**Definition 1. Layer-wise first-order neuron correlation.** For a given neuron  $i$  and neuron  $j$  in layer  $l$ , the layer-wise first-order neuron correlation is given by:

$$A_{ij}^{(l)} = (1/N_{l+1}) \sum_{k=1}^{N_{l+1}} f(W_{i,k}^{(l+1)}) \times f(W_{j,k}^{(l+1)}) \quad (4)$$

The above formula can be expressed as the product of two matrices:

$$A^{(l)} = (1/N_{l+1}) f(W^{(l+1)}) \cdot f(W^{(l+1)})^T \quad (5)$$

where  $\cdot$  represents the matrix multiplication operator.

The layer-wise neuron correlation matrix  $A^{(l)}$  is a symmetric square matrix that models all the pairwise correlations of neurons with respect to their corresponding outgoing weights in layer  $l$ . Each entry  $A_{ij}^{(l)}$  takes a value in the range  $[0, 1)$  and models the correlation between neuron  $i$  and neuron  $j$  in terms of the similarity of their connectivity patterns. The higher the value, the stronger the correlation between the two.

In this setting, two neurons  $i$  and  $j$  from layer  $l$  will be linked and correlated by an intermediate node  $k$  from layer  $l + 1$  if and only if both edges  $f(W_{i,k}^{(l+1)})$  and  $f(W_{j,k}^{(l+1)})$  are non zero, and the relative strength can be estimated by  $f(W_{i,k}^{(l+1)}) \times f(W_{j,k}^{(l+1)})$ , which will be in the range  $[0, 1)$ . Since there are  $N_{l+1}$  neurons in layer  $l + 1$ , where each neuron  $k$  can contribute to such connections, running over all neurons in layer  $l + 1$  we obtain (Equations 4, 5).

### 2.3. The Second-Order Layer-Wise Neuron Correlation

Although the first-order correlation is able to estimate the degree of dependency between each pair of neurons, it may not be sufficient to strictly reflect the degree of grouping or assembly of the neurons. Thus, here we further propose a second-order neuron correlation based on the first-order correlation defined in Equations (4, 5), as:

**Definition 2. Layer-wise second-order neuron correlation.** For a given neuron  $i$  and neuron  $j$  in layer  $l$ , the layer-wise second-order neuron correlation is given by:

$$A_{ij}^{(l)} = (1/N_{l+1}^2) \sum_{k,m} f(W_{i,k}^{(l+1)}) \times f(W_{j,k}^{(l+1)}) \times f(W_{i,m}^{(l+1)}) \times f(W_{j,m}^{(l+1)}) \quad (6)$$

The above formula can be expressed as the product of four matrices:

$$A^{(l)} = (1/N_{l+1}^2) (f(W^{(l+1)}) \cdot f(W^{(l+1)})^T) \odot (f(W^{(l+1)}) \cdot f(W^{(l+1)})^T) \quad (7)$$

where  $\odot$  represents the element-wise multiplication of matrices.

The second-order correlation provides a stricter criterion for relating neurons, as it requires at least two common neighbor nodes from the layer above to have strong connectivity, as compared to the first-order correlation that requires just one common neighbor. Moreover, the second-order neuron correlation is closely related both to graph theory concepts and a neuroscience-inspired learning rule:

**Remark 2. Graph theory and neuroscience interpretation.** Modeling the first-order correlation between two neurons within



the same layer is based on the co-connection to a common neighbor neuron from the layer above, which is closely related to the concepts of clustering coefficient (Watts and Strogatz, 1998) and transitivity (Holland and Leinhardt, 1971) in graph theory. On the other hand, modeling the second-order correlation between two neurons involves two common neighbor neurons in the layer above, which is closely related to calculating the 4-cycle pattern where all 4 possible connections in between are taken into account, as shown in **Figure 1B**. This 4-cycle pattern is linked to the global clustering coefficients of bipartite networks (Robins and Alexander, 2004), where the set of vertices can be decomposed into two disjoint sets such that no two vertices within the same set are adjacent. Similarly, if we consider neurons within one layer as the nodes that belong to one set of the bipartite network between two adjacent layers of the neural networks, forming this 4-cycle will tend to increase the clustering coefficients of the network. Moreover, the second-order correlation is also related to several cognitive neuroscience studies, such as the BIG-ADO learning rule and the principal semantic components of language (Samsonovich et al., 2010; Mainetti and Ascoli, 2015) as well as the notion of discrete neuronal circuits (Pulvermüller and Knoblauch, 2009). **Figure 1A** illustrates a scenario of the BIG-ADO learning rule in BNNs. The blue blobs represents a connection that was formed between two neurons (i.e., a synapse), while the dashed circle between neurons  $j$  and  $m$  represents an Axo-Dendritic Overlap (ADO) (i.e., a potential synapse) between the two neurons. BIG-ADO posits that in order to form a synapse, there must be a potential synapse in place, and the probability of having a potential synapse grows with the second-order correlation. Notably, both of the neuroscience papers cited above relate such a learning mechanism to the formation of cell assemblies in the brain, which parallels our observation of neuronal functional clusters among neurons in DNNs when BEAN was imposed, as shown in **Figures 1C, 6B**.

### 3. EXPERIMENTAL STUDY

Our description of the empirical analysis design and results is organized in the following fashion. In section 3.1, we first characterize the interpretable patterns from the learning outcomes of BEAN regularization on multiple classic image recognition tasks. We then further analyze in section 3.2 how BEAN could benefit the model from learning sparse and efficient neuron connections. Finally, in section 3.3 we study the effect of BEAN regularization on improving the generalizability of the model on several few-shot learning from scratch task simulations. We refer to both distinct BEAN variations, BEAN-1 and BEAN-2, based on the two proposed layer-wise neuron correlation defined by Equations (5, 7) respectively. The value for  $\gamma$  (Equation 3) was set to 1. This paper focuses on examining the effects of the proposed regularization rather than the differences between distinct types of neural network architectures. Hence, we simply adopted several of the most popular neural network architectures for the chosen datasets and did not perform any hyperparameter or system parameter tuning using the test set; in other words, we did not perform any “post selection” (i.e., selectively reporting the model results based on testing set Zheng and Weng, 2016a,b).

All network architectures used in this paper are fully described in their respective cited references, including the specification of their system parameters. The regularization factor of BEAN and other baseline methods were chosen based on the model performance on the validation set. All the experiments were conducted on a 64-bit machine with Intel(R) Xeon(R) W-2155 CPU 3.30 GHz processor and 32 GB memory and an NVIDIA TITAN Xp GPU.

#### 3.1. The Interpretable Patterns of BEAN Regularization

Due to the highly complex computation among numerous layers of neurons in traditional DNNs, it is typically difficult to understand how the network learned what it remembers and the system is more commonly treated as a black-box model (Zhang et al., 2018). Here, to ascertain the effect of BEAN regularization on the interpretability of network dynamics, we analyze the differences in neuronal representation properties of the DNNs with and without BEAN regularization. We conducted experiments on three classic image recognition tasks on the MNIST (LeCun et al., 1998), Fashion-MNIST (Xiao et al., 2017), and CIFAR-10 (Krizhevsky and Hinton, 2009) datasets by starting with three predefined network architectures as listed below:

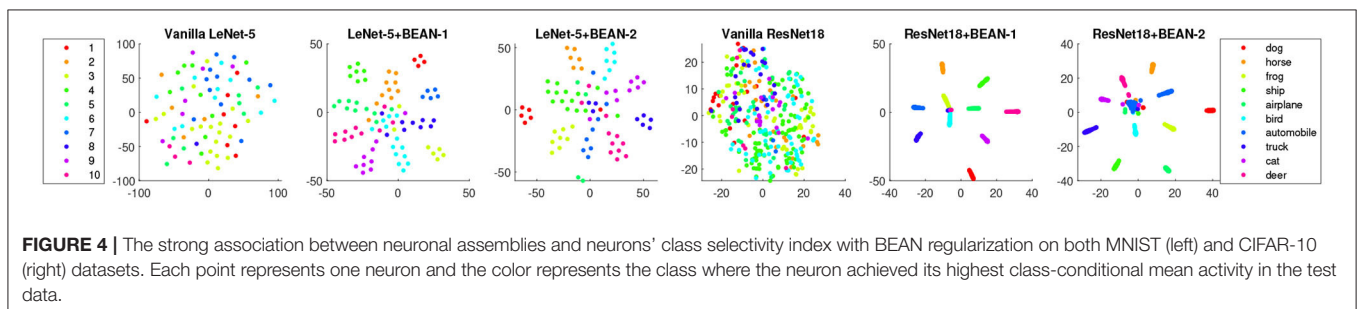
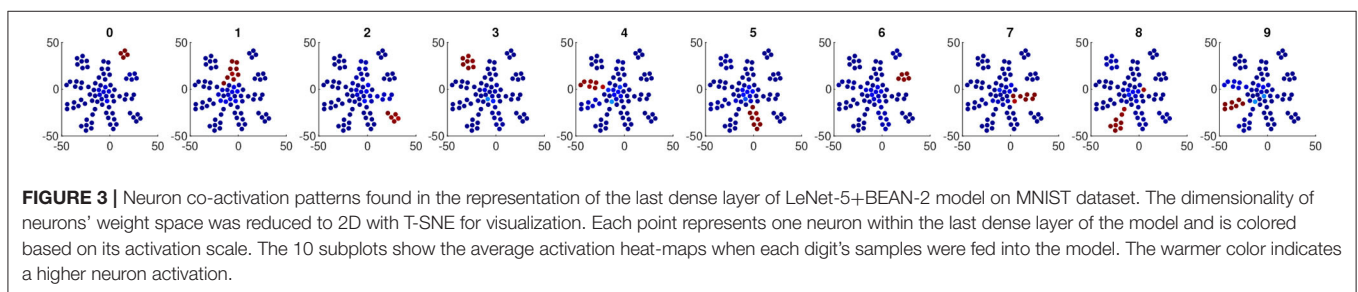
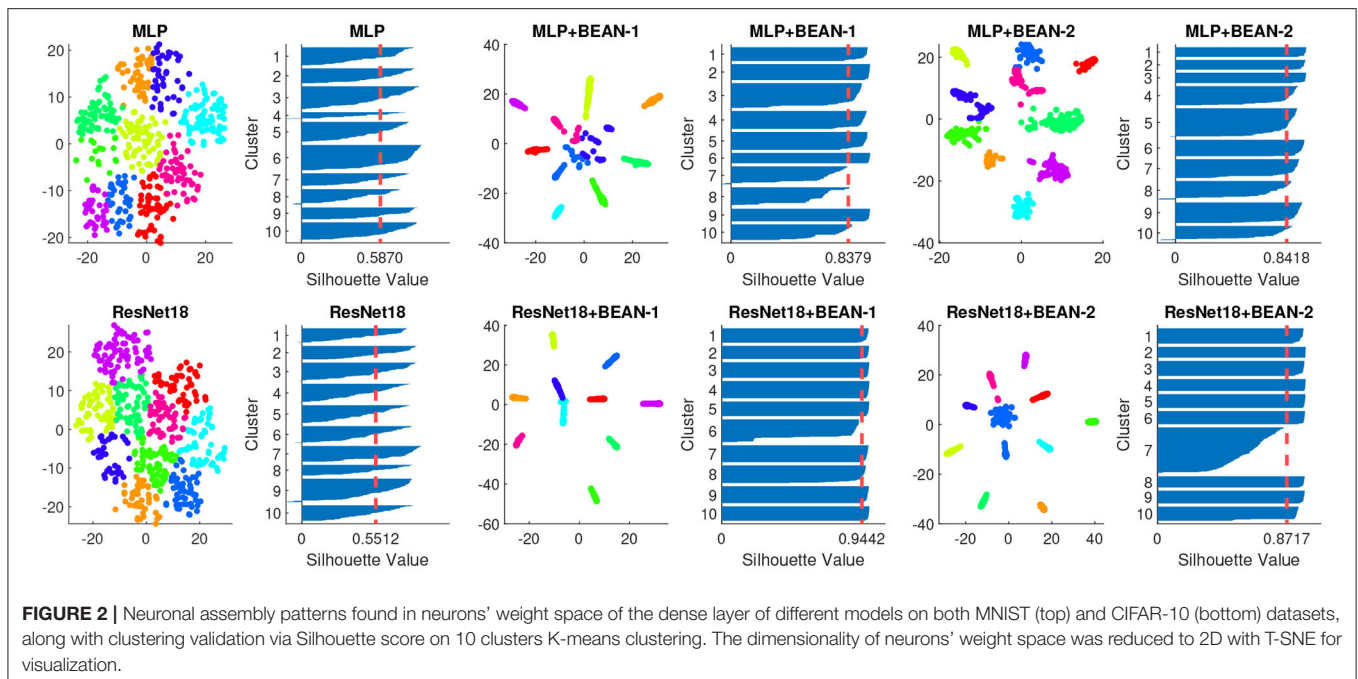
1. An MLP with one hidden layer of 500 neurons with ReLU activation function for MNIST and Fashion-MNIST datasets.
2. A LeNet-5 (LeCun et al., 1998) for MNIST and Fashion-MNIST datasets.
3. ResNet18 (He et al., 2016) for CIFAR-10 dataset.

The Adam optimizer (Kingma and Ba, 2014) was used with a learning rate of 0.0005 and a batch size of 100 for model training until train loss convergence was achieved; BEAN was applied to all the dense layers of each model.

##### 3.1.1. Biological Plausibility of the Learned Neuronal Assemblies

By analyzing the neurons' connectivity patterns based on their outgoing weights, we discovered neuronal assemblies in dense layers where BEAN regularization was enforced. Specifically, for both datasets, we found that the neuronal assemblies at the last dense layer could be best described by 10 clusters with K-means clustering (MacQueen, 1967) validated by Silhouette analysis (Rousseeuw, 1987). Silhouette analysis is a widely-used method for interpretation and validation of consistency within clusters of data. The technique provides a succinct graphical representation of how well each object has been classified. As shown in **Figure 2**, we visualized the K-means clustering results in neurons' weight space of the dense layer on both MNIST (top) and CIFAR-10 (bottom) datasets. Each data point in the figure indicates one single neuron and the color indicates its cluster assignment by the clustering algorithm. The Silhouette value is further used to assess the quality of the clustering assignment: high Silhouette values support the existence of clear clusters in the data points, which here correspond to neural assembly patterns among neurons.

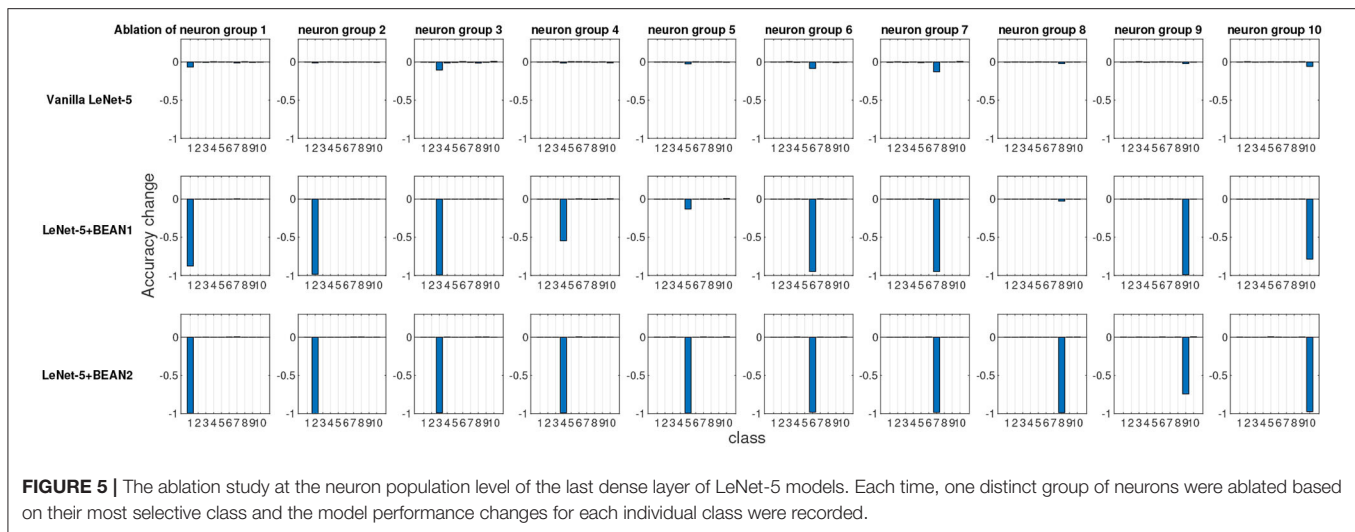
Both BEAN-1 and BEAN-2 could enforce neuronal assemblies for various models on several datasets, yielding Silhouette indices



around 0.9, which indicates strong clustering patterns among neurons in dense layers where BEAN regularization was applied. On the other hand, training conventional DNN models with the same architectures could only yield Silhouette indices near 0.5, which indicates no clear clustering patterns in conventional dense layers of deep neuronal networks.

Moreover, we found co-activation behavior of neurons within each neuronal assembly, which is both interpretable and biologically plausible. **Figure 3** shows the visualization of neuron co-activation patterns found in the last dense

layer of LeNet-5+BEAN-2 model on MNIST dataset. For the samples of each specific class, only those neurons in the specific neuron group that is associated with that digit class have high activation while all the other neurons remain silent. This strong correlation between each unique assembly and each unique class concept allows straightforward interpretation of the neuron populations in the dense layers. From the neuroscience perspective, those co-activation patterns and the association between high-level concepts and neuron groups may reflect similar co-firing patterns



observed in biological neural systems (Peyrache et al., 2010) and underscore the strong association between neuronal assembly and concepts (Tononi and Sporns, 2003) in biological neural networks.

We also found a strong correlation between neuronal assembly and class selectivity indices. Selectivity index was originally proposed and used in systems neuroscience (De Valois et al., 1982; Freedman and Assad, 2006). Recently, machine learning researchers also studied unit class selectivity (Morcos et al., 2018; Zhou et al., 2018) as a metric for interpreting the behaviors of single units in deep neural networks. Mathematically, it is calculated as:  $selectivity = (\mu_{max} - \mu_{-max}) / (\mu_{max} + \mu_{-max})$ , where  $\mu_{max}$  represents the highest class-conditional mean activity and  $\mu_{-max}$  represents the mean activity across all other classes.

To better visualize how high-level concepts are associated with the learned neuron assemblies, we further labeled each neuron with the class in which it achieved its highest class-conditional mean activity  $\mu_{max}$  in the test data. **Figure 4** shows the results for the last dense layer of the models trained with both datasets. We found that the neuronal assembly could be well described based on selectivity. The strong association between neuronal assemblies and neurons' selectivity index further demonstrated the biological plausibility of the learning outcomes of BEAN regularization. Moreover, the strong neuron activation patterns toward each individual high-level concepts or classes could in principle enable one to better understand what each individual neuron has learned to represent. However, more relevant to and consistent with our regularization, these selective activation patterns reveal how a group of neurons (i.e., neuronal assembly) together capture the whole picture of each high-level concept, such as the 'bird' class in CIFAR-10 as shown in **Figure 4**.

In this subsection, we have demonstrated the promising effect of the proposed BEAN regularization on forming the neural assembly patterns among the neurons in the last layer of the network and their correspondence with biological neural networks. Although the effect of BEAN regularization is not yet

clear on the lower layers of the networks, it will be interesting in the future to explore additional relations between computational function and the architecture of earlier processing stations in biological neural systems.

### 3.1.2. Quantitative Analysis of Interpretability

Experimental neuropsychologists commonly use an ablation protocol when studying neural function, whereas parts of the brain are removed to investigate the cognitive effects. Similar ablation studies have also been adapted for interpreting deep neural networks, such as understanding which layers or units are critical for model performance (Girshick et al., 2014; Morcos et al., 2018; Zhou et al., 2018).

To quantitatively evaluate and compare interpretability, we performed an ablation study at the neuron population level, each time ablating one distinct group of neurons and recording the consequent model performance changes for each class. As shown in **Figure 4**, we identified neuron groups via class selectivity and performed neuron population ablation accordingly. **Figure 5** shows the results of all 10 ablation runs for each class in MNIST dataset. As also reported by Morcos et al. (2018), for conventional deep neural nets, there is indeed no clear association between neuron's selectivity and importance to the overall model performance, as revealed by neuron population ablation. However, when BEAN regularization was utilized during training, such association clearly emerged, especially for BEAN-2. This is because BEAN-2 could enforce neurons to form stricter neuron correlations than BEAN-1 with the second-order correlation, enabling groups of neurons to represent more compact and disentangled concepts, such as handwritten digits. This discovery further demonstrated the interpretability and concept level representation in each neuronal assembly learned by applying BEAN regularization. Such compact and interpretable structure of concept-level information encoding could also benefit the field of disentanglement representation learning (Bengio et al., 2013).

**TABLE 1** | Efficient model learning experiments on MNIST and Fashion-MNIST datasets.

Dataset	Measure	Vanilla	$\ell_1$ -norm	Group sparsity	Exclusive sparsity	BEAN-1	BEAN-2
MNIST	Accuracy	0.9812	<i>0.9835</i>	0.9813	0.9824	<b>0.9842</b>	0.9823
	FLOPs	1	0.8106	0.6098	0.4248	<i>0.2212</i>	<b>0.1320</b>
	Parameter	1	0.2921	<i>0.0982</i>	0.1375	0.1496	<b>0.0730</b>
Fashion-MNIST	Accuracy	<b>0.8986</b>	0.8924	0.8925	0.8930	<i>0.8960</i>	0.8916
	FLOPs	1	0.8011	0.5384	0.5320	<i>0.2913</i>	<b>0.1622</b>
	Parameter	1	0.4357	<i>0.1378</i>	0.2257	0.2592	<b>0.1259</b>

The FLOPs and effective parameters (i.e., number of non-zero parameters) are normalized by the value of vanilla model. Performance is averaged over 20 runs. The best and second-best results are highlighted in boldface and italic font, respectively.

**TABLE 2** | Few-shot learning from scratch experiments on the MNIST (left), Fashion-MNIST (middle), and CIFAR-10 (right) datasets.

Dataset	MNIST				Fashion-MNIST				CIFAR-10			
	1-shot	5-shot	10-shot	20-shot	1-shot	5-shot	10-shot	20-shot	1-shot	5-shot	10-shot	20-shot
Model	1-shot	5-shot	10-shot	20-shot	1-shot	5-shot	10-shot	20-shot	1-shot	5-shot	10-shot	20-shot
Vanilla	38.63	70.21	78.97	86.68	39.32	59.02	64.50	70.23	15.60	18.49	22.45	26.39
Dropout	40.13	72.45	82.04	89.22	40.78	60.04	65.40	71.83	15.10	18.85	22.73	26.01
Weight decay	39.51	71.76	82.87	90.15	41.31	61.98	67.25	71.88	15.47	19.17	23.74	26.77
$\ell_1$ -norm	40.96	74.35	81.17	90.68	41.26	62.18	67.30	70.85	15.64	18.95	23.16	26.99
Imitation networks	44.10	70.40	80.00	86.70	44.80	62.10	68.00	72.50	-	-	-	-
BEAN-cos	<i>54.05</i>	80.16	86.28	92.22	42.48	65.49	68.97	74.20	18.23	<i>21.45</i>	24.66	28.74
BEAN-1	<b>54.79</b>	<b>83.42</b>	<i>87.51</i>	92.79	<b>50.57</b>	<b>66.95</b>	<i>69.21</i>	<i>74.25</i>	<b>19.39</b>	<b>21.92</b>	<i>24.81</i>	28.95
BEAN-2	53.75	<i>80.76</i>	<b>88.08</b>	<b>92.97</b>	<i>49.94</i>	65.98	<b>70.21</b>	<b>75.06</b>	19.28	21.28	<b>25.04</b>	<b>29.23</b>

Performance is averaged over 20 simulations of randomly sampled training data from the original training base. The best and second-best results for each few-shot learning setting are highlighted in boldface and italic font, respectively.

**TABLE 3** | Statistic of data values test set error rate-validation set error rate on 10-shot learning on the MNIST dataset from 20 random runs.

Model / Metric	Max	75%-rank	50%-rank	Mean	25%-rank	Min
Vanilla	0.06%	0.01%	-0.06%	-0.30%	-0.74%	-0.81%
BEAN-1( $\alpha = 1$ )	-0.04%	-0.20%	-0.62%	-0.58%	-0.93%	-1.13%
BEAN-2( $\alpha = 100$ )	0.10%	-0.02%	-0.42%	-0.48%	-0.97%	-1.35%

Other  $n$ -shot learning settings follow the same trend.

### 3.2. Learning Sparse and Efficient Networks

To evaluate the effect of BEAN regularization on learning sparse and efficient networks, we conducted experiments on two real-world benchmark datasets, i.e., the MNIST (LeCun et al., 1998) and Fashion-MNIST (Xiao et al., 2017) datasets. We compared BEAN with several state-of-the-art regularization methods that could enforce sparse connection of the network, including  $\ell_1$ -norm, group sparsity based on  $\ell_{2,1}$ -norm (Yuan and Lin, 2006; Alvarez and Salzmann, 2016), and exclusive sparsity based on  $\ell_{1,2}$ -norm (Zhou et al., 2010; Kong et al., 2014). Notable studies also investigated the combination of the sparsity terms listed above, such as combining group sparsity and  $\ell_1$ -norm (Scardapane et al., 2017), and combining group and exclusive sparsity (Yoon and Hwang, 2017). The combinatorial study is outside the scope of this work, as our focus is on showing and comparing the effectiveness of the single regularization term to the network. To keep the comparison fair and accurate, we use

the same base network architecture for all regularization methods tested in this experiment, which is a predefined fully connected neural network with three hidden layers, 500 neurons per layer, and ReLU as the neuron activation function. The regularization methods are applied to all layers of the network, except the bias term. The regularization co-efficients are selected through a grid search varying from  $10^{-5}$  to  $10^3$  based on the model performance on the validation set, as shown in Algorithm 1. To obtain a more reliable and fair result, we ran a total of 20 random weight initializations for every network architecture studied and reported the overall average performance of all 20 results as the final model performance of each architecture.

To quantitatively measure the performance of various sparse regularization techniques, we used three evaluation metrics, including the prediction accuracy on test data (i.e., measured by the number of correct predictions divided by the total number of samples in test data), the ratio of parameters used in the network (i.e., total number of non-zero weights divided by the



---

**Algorithm 1:** The pseudo code for searching for the best  $\alpha$  value in BEAN
 

---

```

func hyperparameter_tuner(training_data, validation_data, alpha_list =
[0.001, 0.01, 0.1, 1, 10, 100]):
  hp_perf = []
  % train and evaluate on all hyper-parameter settings
  foreach  $\alpha$  in alpha_list:
    m = train_model(training_data, alpha)
    validation_results = eval_model(m, validation_data)
    hp_perf.append(validation_results)
  % find the best alpha on validation set
  best_alpha = alpha_list[max_index(hp_perf)]
  return best_alpha

```

---

total number of weights in the networks after training), and the corresponding number of floating point operations (FLOPs). A higher accuracy means that the model can train a better network for the classification tasks. A lower FLOP indicates that the network needs fewer computational operations per forwarding pass, which reflects computation efficiency. Similarly, a lower parameter usage indicates the network requires less memory usage, which reflects memory efficiency.

The results are shown in **Table 1**. For each evaluation metric, the best and second-best results are highlighted in boldface and italic font, respectively. As can be seen, both BEAN-1 and BEAN-2 can achieve high memory and computational efficiency without sacrificing network performance for the classification tasks. Specifically, BEAN-2 achieved the best memory and computational efficiency, out-performing baseline models by 25–75% on memory efficiency and 69–84% on computational efficiency on the MNIST dataset, and by 9–71% on memory efficiency and 69–80% on computational efficiency on the Fashion-MNIST dataset. BEAN-1 also achieved a good trade-off between model performance and efficiency, being the second-best on computational efficiency and the best on model performance on both the MNIST and Fashion-MNIST datasets. Comparing with BEAN-2, BEAN-1 leans more toward the model performance side in such a trade-off. This is because the first-order correlation used in BEAN-1 is less restrictive than a higher-order correlation in BEAN-2, as only one support neuron in the layer above is enough to build up a strong correlation. Thus, in practice, using a higher-order correlation might be promising when the objective is to learn a more efficient model.

Interestingly, BEAN regularization seemed to advance the state-of-the-art by an even more significant margin in terms of computational efficiency. In fact, BEAN regularization reduces the number of FLOPs needed for the network by automatically “pruning” a substantial proportion of neurons in the hidden layers (whereas a neuron is considered pruned if either all incoming or all outgoing weights are zero), due to the penalization of connections between neurons that encode divergent information. Although group sparsity and exclusive sparsity are designed to achieve a similar objective for obtaining neuron-level sparsity, they are less effective than BEAN regularization. This is due to the fact that BEAN takes into consideration not only the correlations between neurons via their

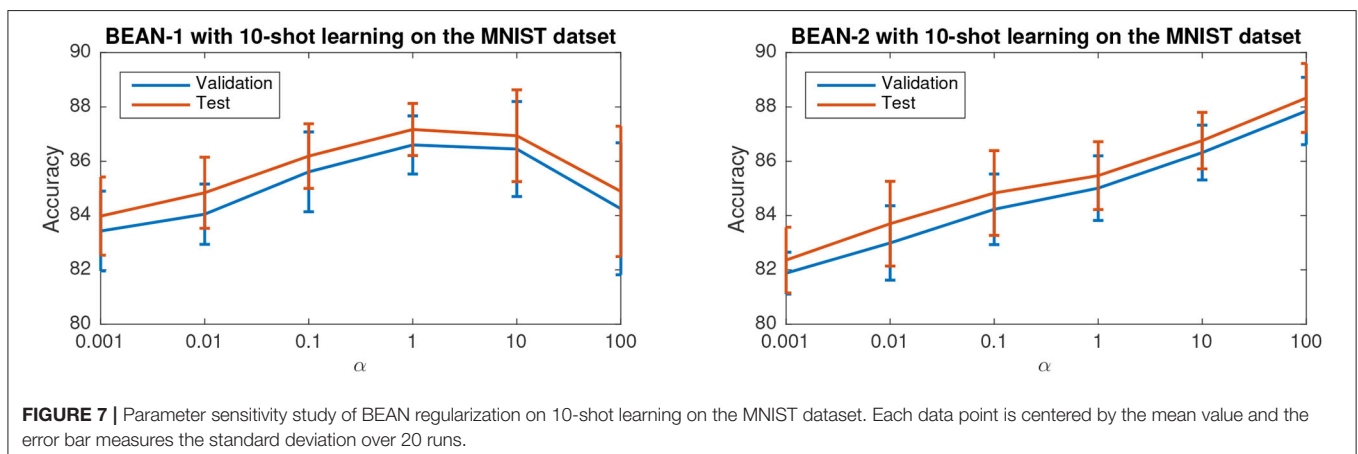
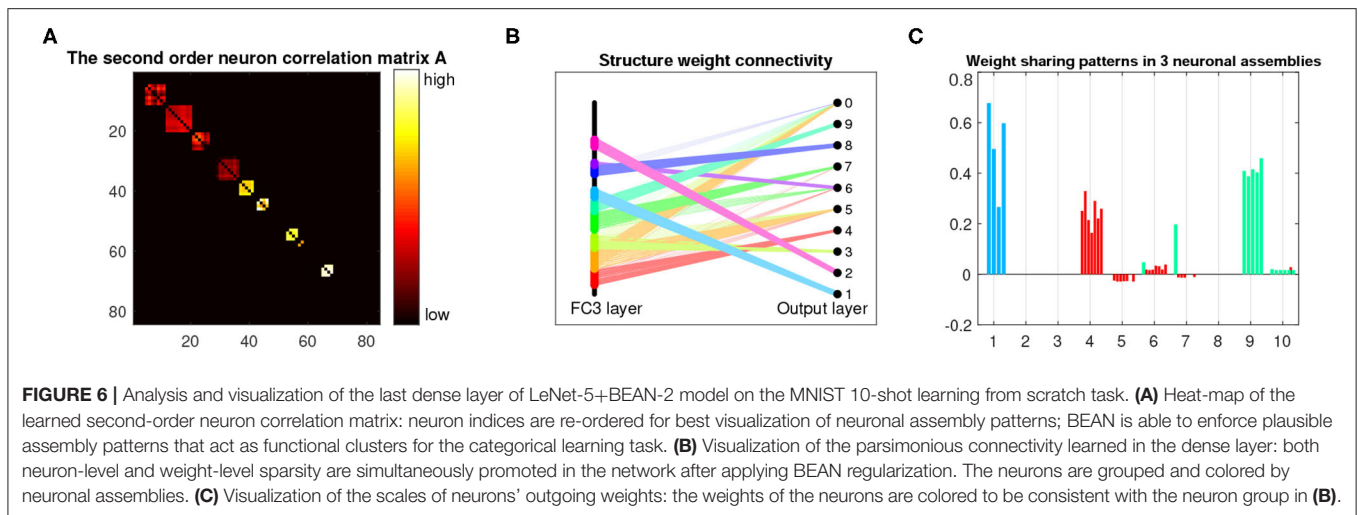
connection patterns but also the consistency of those correlations with their activation patterns.

We have shown in **Table 1** that the proposed BEAN regularization can effectively make the connection sparser in the dense layers of the artificial neural networks. In general, this “sparsifying” effect can be beneficial for any models with at least one dense layer in the network architecture. Most modern deep neural networks (such as VGG Simonyan and Zisserman, 2014 and ImageNet Russakovsky et al., 2015) can enjoy this sparsity benefit, as the dense layers typically contribute to the majority of the model parameters (Cheng et al., 2015).

### 3.3. Toward Few-Shot Learning From Scratch With BEAN Regularization

In an attempt to test the influence of BEAN regularization on the generalizability of DNNs in the scenarios where the training samples are extremely limited, we conducted a *few-shot learning from scratch* task, i.e., without the help of any additional side tasks and pre-trained models (Kimura et al., 2018). Notice that in the few-shot learning setting, the model typically requires an iterative learning process over the sample set. In other words, for each individual few-shot learning experiment, only a few image samples per digit are randomly selected to form the training set. The model then iteratively learns from the selected image samples until convergence is achieved. So far, this kind of learning task has rarely been explored due to the difficulty of the problem setup as compared to other conventional few-shot learning tasks where additional data or knowledge could be accessed. Currently, only (Kimura et al., 2018) carried out a preliminary exploration with their proposed Imitation Networks model. We conducted several simulations of the *few-shot learning from scratch* task on the MNIST (LeCun et al., 1998), Fashion-MNIST (Xiao et al., 2017), and CIFAR-10 (Krizhevsky and Hinton, 2009) datasets. Besides Kimura’s Imitation Networks, we also compared BEAN with other conventional regularization techniques commonly used in the deep learning literature. Specifically, we compared dropout (Srivastava et al., 2014), weight decay (Krogh and Hertz, 1992), and  $\ell_1$ -norm. Similarly to the description of section 3.2, we kept the comparison fair and accurate by using a predefined network architecture, namely LeNet-5 (LeCun et al., 1998), as the base network architecture for all regularization methods studied in this experiment. The regularization terms were applied to all three dense layers of the base LeNet-5 network. Once again, the hyperparameter of each regularization along with all other system parameters were selected through a grid search and based on the best performance on a predefined 10k validation set sampled from the original training base and completely distinct from the training samples used in the few-shot learning tasks and the testing set.

**Table 2** shows model performance on several *few-shot learning from scratch* experiments on the MNIST, Fashion-MNIST, and CIFAR-10 datasets. Performance is averaged over 20 experiments of randomly sampled training data from the original training base. The best and second-best results for each few-shot learning settings are highlighted in boldface and italic font, respectively. As can be seen, the proposed BEAN regularization advanced



the state-of-the-art by a significant margin on all four *few-shot learning from scratch* tasks tested among all three datasets. Moreover, BEAN advanced the performance more significantly when training samples were more limited. For instance, BEAN outperformed all comparison methods by 24–42%, 13–29%, and 24–28% on 1-shot learning tasks on the MNIST, Fashion-MNIST, and CIFAR-10 datasets, respectively. This observation demonstrates the promising effect of BEAN regularization on improving the generalizability of the neural nets when the training samples are extremely limited. Another interesting observation is that BEAN-1 in general performed the best with extremely limited training samples, such as the 1-shot and 5-shot learning tasks, while BEAN-2 regularization in general performed the best with slightly more training samples, such as the 10-shot and 20-shot learning tasks. The reason behind this observation might be related to the more stringent higher-order correlation, which requires more common neighbor neurons that appear to have strong connections with both neurons. Thus, a modestly increased availability of sample observations could enable BEAN-2 to form more effective neuronal assemblies, further improving the model performance.

Furthermore, we studied an additional variant for BEAN, i.e., BEAN-cos, which calculates the layer-wise neuron correlation via cosine similarity between the downstream weights of two neurons. As shown in **Table 2**, we found that BEAN-cos can still yield good performance and beat other existing regularization methods, and getting competitive results as compared with BEAN-1 and BEAN-2. However, it is inferior to BEAN-1 in 1-shot and 5-shot settings, and inferior to BEAN-2 in 10-shot and 20-shot settings. This is because BEAN-cos is unable to handle the order of correlation between neurons, as using cosine similarity requires us to treat the out-going weights of a neuron as a whole (vector) to compute the pairwise similarity between neurons. Thus, doing this will lose the ability to calculate higher-order correlation (such as the second-order correlation), and consequentially lose the good interpretation from graph theory and neuroscience (as described in Remark 2).

To better understand why BEAN regularization could help the seemingly over-parameterized model generalize well on a small sample set, we further analyzed the learned hidden representation of the dense layers where

BEAN regularization was employed. We found that BEAN helped the model gain better generalization power in two aspects: (1) by automatic sparse and structured connectivity learning and (2) by weak parameter sharing among neurons within each neuronal assembly. Both aspects enhanced the dense layers to promote efficient and parsimonious connections, which consequently prevented the model from over-fitting with a small training sample size.

**Figure 6** shows the learned parameters of the last dense layer of LeNet-5+BEAN2 on the MNIST 10-shot learning task. As shown in **Figure 6B**, instead of using all possible weights in the dense layer, BEAN caused the model to parsimoniously leverage the weights and even the neurons, yielding a bio-plausible sparse and structured connectivity pattern. This is because the learned neuron correlation helped the model disentangle the co-connections between neurons from different assemblies, as shown in **Figure 6A**. Additionally, BEAN enhanced parameter sharing among neurons within each assembly, as demonstrated in **Figure 6C**. For instance, neurons in the red-colored assembly all had high positive weights toward class 4, meaning that this group of neurons was helping the model identify Digit 4. Similarly, neurons in the green-colored assembly were trying to distinguish between Digits 9 and 7. Such automatic weak parameter sharing not only helped prevent the model from over-fitting but also enabled an intuitive interpretation of the behavior of the system as a whole from a higher modularity level.

### 3.3.1. Parameter Sensitivity Study

There are two hyperparameters in the proposed BEAN regularization: (1)  $\alpha$ , which balances between the regularization loss and DNN training loss, and (2)  $\gamma$ , which controls the curvature of the hyperbolic tangent function as shown in Equation (3). As already mentioned in the first paragraph of section 3,  $\gamma$  was set to 1 for all experiments. Thus, the only parameter we need to study is  $\alpha$ .

**Figure 7** shows the accuracy of the model versus  $\alpha$  on the few-shot learning setting on the MNIST dataset. Only the results for the 10-shot learning task are shown due to space limitations. By varying  $\alpha$  across the range from 0.001 to 100, the best performance is obtained when  $\alpha = 1$  for BEAN-1 and  $\alpha = 100$  for BEAN-2. Specifically, for BEAN-1, we can see a clear trend where the model performance drops when  $\alpha$  is too small or too big. Furthermore, the results show that the performance of the validation set is well aligned with the model performance on the test set, as also shown in the statistical analysis results in **Table 3**. This demonstrates the superior generalizability of the model when applying BEAN regularization. Notably, although in **Figure 7** we accessed the model performance on multiple settings of  $\alpha$ , we did not use any of the results on the test set to choose any parameters of the model, i.e., no post-selection was performed. We believe post-selection should be completely avoided and it can cause the test set to lose its power to test the model's generalizability to future unseen data.

## 4. CONCLUSION

In this paper, we propose a novel Biologically Enhanced Artificial Neuronal assembly (BEAN) regularization to model neuronal correlations and dependencies inspired by cell assembly theory from neuroscience. We show that BEAN can promote jointly sparse and efficient encoding of rich semantic correlation among neurons in DNNs similar to connection patterns in BNNs. Experimental results show that BEAN enables the formations of interpretable neuronal functional clusters and consequently promotes a sparse, memory/computation-efficient network without loss of model performance. Moreover, our few-shot learning experiments demonstrated that BEAN could also enhance the generalizability of the model when training samples are extremely limited. Our regularization method has demonstrated its capability in enhancing the modularity of the representations of neurons for image semantic meanings such as digits, animals, and objects on image datasets. While the generality of the approach introduced here is at this time evaluated on MNIST and CIFAR datasets, future studies might consider additional experiments on other datasets such as texts or graphs to demonstrate the broader effectiveness of the proposed method. Another direction to further enhance the model might be to include separate excitatory and inhibitory nodes, as in BNNs, which would allow implementation of specific microcircuit computational motifs (Ascoli and Atkeson, 2005). Furthermore, since there are other choices for defining the affinity matrix between neurons in a certain layer based on their downstream weights, answering the question about "what is the best way to compute affinity matrix" can be an interesting direction to be more comprehensively studied in future works.

## DATA AVAILABILITY STATEMENT

Publicly available datasets were analyzed in this study. This data can be found here: MNIST: <http://yann.lecun.com/exdb/mnist/>; Fashion-MNIST: <https://github.com/zalandoresearch/fashion-mnist>; CIFAR-10: <https://www.cs.toronto.edu/~kriz/cifar.html>.

## AUTHOR CONTRIBUTIONS

YG and LZ conceived of the presented idea. YG developed the model and performed the experiment simulations. GA verified the computational and neuroscience plausibility of the proposed method. YG initialized the manuscript and all authors discussed the results and contributed to the final manuscript.

## FUNDING

This work is supported by the National Institutes of Health grant (NS39600), the National Science Foundation grant: #1755850, #1841520, #1907805, Jeffress Trust Award, NVIDIA GPU Grant, and the Design Knowledge Company (subcontract number: 10827.002.120.04). This manuscript has been released as a preprint at arXiv: 1909.13698 (Gao et al., 2019).

## REFERENCES

- Alvarez, J. M., and Salzmann, M. (2016). "Learning the number of neurons in deep networks," in *Advances in Neural Information Processing Systems* (Barcelona), 2270–2278.
- Ascoli, G. A., and Atkeson, J. C. (2005). Incorporating anatomically realistic cellular-level connectivity in neural network models of the rat hippocampus. *Biosystems* 79, 173–181. doi: 10.1016/j.biosystems.2004.09.024
- Bengio, Y., Courville, A., and Vincent, P. (2013). Representation learning: a review and new perspectives. *IEEE Trans. Pattern Anal. Mach. Intel.* 35, 1798–1828. doi: 10.1109/TPAMI.2013.50
- Braitenberg, V. (1978). "Cell assemblies in the cerebral cortex," in *Theoretical Approaches to Complex Systems*, (Berlin; Heidelberg: Springer), 171–188.
- Buzsáki, G. (2010). Neural syntax: cell assemblies, synapses, and readers. *Neuron* 68, 362–385. doi: 10.1016/j.neuron.2010.09.023
- Cheng, Y., Yu, F. X., Feris, R. S., Kumar, S., Choudhary, A., and Chang, S.-F. (2015). "An exploration of parameter redundancy in deep networks with circulant projections," in *Proceedings of the IEEE International Conference on Computer Vision* (Santiago), 2857–2865.
- Collobert, R., and Weston, J. (2008). "A unified architecture for natural language processing: deep neural networks with multitask learning," in *Proceedings of the 25th International Conference on Machine Learning* (Helsinki: ACM), 160–167.
- De Valois, R. L., Yund, E. W., and Hepler, N. (1982). The orientation and direction selectivity of cells in macaque visual cortex. *Vis. Res.* 22, 531–544. doi: 10.1016/0042-6989(82)90112-2
- Finn, C., Abbeel, P., and Levine, S. (2017). "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proceedings of the 34th International Conference on Machine Learning-Vol. 70* (Sydney, NSW: JMLR.org), 1126–1135.
- Freedman, D. J., and Assad, J. A. (2006). Experience-dependent representation of visual categories in parietal cortex. *Nature* 443:85. doi: 10.1038/nature05078
- Gao, Y., Ascoli, G., and Zhao, L. (2019). Bean: interpretable representation learning with biologically-enhanced artificial neuronal assembly regularization. *arXiv preprint arXiv:1909.13698*.
- Girshick, R., Donahue, J., Darrell, T., and Malik, J. (2014). "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Columbus, OH), 580–587.
- Granovetter, M. S. (1973). The strength of weak ties. *Am. J. Sociol.* 78, 1360–1380.
- He, K., Zhang, X., Ren, S., and Sun, J. (2016). "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Las Vegas, NV), 770–778.
- Hebb, D. O. (1949). *The Organization of Behavior. A Neuropsychological Theory*. New York, NY: Wiley.
- Holland, P. W., and Leinhardt, S. (1971). Transitivity in structural models of small groups. *Compar. Group Stud.* 2, 107–124. doi: 10.1177/104649647100200201
- Kimura, A., Ghahramani, Z., Takeuchi, K., Iwata, T., and Ueda, N. (2018). Few-shot learning of neural networks from scratch by pseudo example optimization. *arXiv preprint arXiv:1802.03039*.
- Kingma, D. P., and Ba, J. (2014). Adam: a method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Kong, D., Fujimaki, R., Liu, J., Nie, F., and Ding, C. (2014). "Exclusive feature learning on arbitrary structures via  $\ell_{1,2}$ -norm," in *Advances in Neural Information Processing Systems* (Montreal, QC), 1655–1663.
- Krizhevsky, A., and Hinton, G. (2009). *Learning Multiple Layers of Features From Tiny Images*. Technical report, Citeseer.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems* (Stateline, NV), 1097–1105.
- Krogh, A., and Hertz, J. A. (1992). "A simple weight decay can improve generalization," in *Advances in Neural Information Processing Systems* (Denver, CO), 950–957.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature* 521:436. doi: 10.1038/nature14539
- LeCun, Y., Bottou, L., Bengio, Y., and Haffner, P. (1998). Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 2278–2324. doi: 10.1109/5.726791
- Maaten, L. V. d., and Hinton, G. (2008). Visualizing data using t-sne. *J. Mach. Learn. Res.* 9, 2579–2605.
- MacQueen, J. (1967). "Some methods for classification and analysis of multivariate observations," in *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability Vol. 1*, (Oakland, CA, USA), 281–297.
- Mainetti, M., and Ascoli, G. A. (2015). A neural mechanism for background information-gated learning based on axonal-dendritic overlaps. *PLoS Comput. Biol.* 11:e1004155. doi: 10.1371/journal.pcbi.1004155
- Morcos, A. S., Barrett, D. G., Rabinowitz, N. C., and Botvinick, M. (2018). On the importance of single directions for generalization. *arXiv preprint arXiv:1803.06959*.
- Peyrache, A., Benchenane, K., Khamassi, M., Wiener, S. I., and Battaglia, F. P. (2010). Principal component analysis of ensemble recordings reveals cell assemblies at high temporal resolution. *J. Comput. Neurosci.* 29, 309–325. doi: 10.1007/s10827-009-0154-6
- Pulvermüller, F., and Knoblauch, A. (2009). Discrete combinatorial circuits emerging in neural networks: a mechanism for rules of grammar in the human brain? *Neural Networks* 22, 161–172. doi: 10.1016/j.neunet.2009.01.009
- Rees, C. L., Moradi, K., and Ascoli, G. A. (2017). Weighing the evidence in peters' rule: does neuronal morphology predict connectivity? *Trends Neurosci.* 40, 63–71. doi: 10.1016/j.tins.2016.11.007
- Rivera-Alba, M., Peng, H., de Polavieja, G. G., and Chklovskii, D. B. (2014). Wiring economy can account for cell body placement across species and brain areas. *Curr. Biol.* 24, R109–R110. doi: 10.1016/j.cub.2013.12.012
- Robins, G., and Alexander, M. (2004). Small worlds among interlocking directors: network structure and distance in bipartite graphs. *Comput. Math. Organ. Theory* 10, 69–94. doi: 10.1023/B:CMOT.0000032580.12184.c0
- Ropireddy, D., Scorcioni, R., Lasher, B., Buzsáki, G., and Ascoli, G. A. (2011). Axonal morphometry of hippocampal pyramidal neurons semi-automatically reconstructed after *in vivo* labeling in different CA3 locations. *Brain Struct. Funct.* 216, 1–15. doi: 10.1007/s00429-010-0291-8
- Rousseeuw, P. J. (1987). Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *J. Comput. Appl. Math.* 20, 53–65. doi: 10.1016/0377-0427(87)90125-7
- Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1988). Learning representations by back-propagating errors. *Cogn. Model.* 5:1.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., et al. (2015). Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis.* 115, 211–252. doi: 10.1007/s11263-015-0816-y
- Sabour, S., Frosst, N., and Hinton, G. E. (2017). "Dynamic routing between capsules," in *Advances in Neural Information Processing Systems* (Long Beach, CA), 3856–3866.
- Samsonovich, A. V., Goldin, R. F., and Ascoli, G. A. (2010). Toward a semantic general theory of everything. *Complexity* 15, 12–18. doi: 10.1002/cpl.x.20293
- Scardapane, S., Comminiello, D., Hussain, A., and Uncini, A. (2017). Group sparse regularization for deep neural networks. *Neurocomputing* 241, 81–89. doi: 10.1016/j.neucom.2017.02.029
- Simonyan, K., and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Snell, J., Swersky, K., and Zemel, R. (2017). "Prototypical networks for few-shot learning," in *Advances in Neural Information Processing Systems*, 4077–4087.
- Socher, R., Ganjoo, M., Manning, C. D., and Ng, A. (2013). "Zero-shot learning through cross-modal transfer," in *Advances in Neural Information Processing Systems* (Stateline, NV), 935–943.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* 15, 1929–1958.
- Tononi, G., and Sporns, O. (2003). Measuring information integration. *BMC Neurosci.* 4:31. doi: 10.1186/1471-2202-4-31
- Watts, D. J., and Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature* 393:440. doi: 10.1038/30918
- Xian, Y., Lampert, C. H., Schiele, B., and Akata, Z. (2018). Zero-shot learning—a comprehensive evaluation of the good, the bad and the ugly. *IEEE Trans. Pattern Anal. Mach. Intel.* 41, 2251–2265. doi: 10.1109/CVPR.2017.328



- Xiao, H., Rasul, K., and Vollgraf, R. (2017). Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*.
- Yoon, J. and Hwang, S. J. (2017). “Combined group and exclusive sparsity for deep neural networks,” in *Proceedings of the 34th International Conference on Machine Learning-Vol. 70* (JMLR.org), 3958–3966.
- Young, T., Hazarika, D., Poria, S., and Cambria, E. (2018). Recent trends in deep learning based natural language processing. *IEEE Comput. Intel. Mag.* 13, 55–75. doi: 10.1109/MCI.2018.2840738
- Yuan, M., and Lin, Y. (2006). Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc.* 68, 49–67. doi: 10.1111/j.1467-9868.2005.00532.x
- Zhang, Q., Nian Wu, Y., and Zhu, S.-C. (2018). “Interpretable convolutional neural networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (Salt Lake City, UT), 8827–8836.
- Zhang, Q.-S., and Zhu, S.-C. (2018). Visual interpretability for deep learning: a survey. *Front. Inform. Technol. Electron. Eng.* 19, 27–39. doi: 10.1631/FITEE.1700808
- Zheng, Z., and Weng, J. (2016a). “Challenges in visual parking and how a developmental network approaches the problem,” in *2016 International Joint Conference on Neural Networks (IJCNN)* (Vancouver, BC: IEEE), 4593–4600.
- Zheng, Z., and Weng, J. (2016b). “Mobile device based outdoor navigation with on-line learning neural network: a comparison with convolutional neural network,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops* (Las Vegas, NV), 11–18.
- Zhou, B., Sun, Y., Bau, D., and Torralba, A. (2018). Revisiting the importance of individual units in cnns via ablation. *arXiv preprint arXiv:1806.02891*.
- Zhou, Y., Jin, R., and Hoi, S. C.-H. (2010). “Exclusive lasso for multi-task feature selection,” in *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics* (Sardinia), 988–995.

**Conflict of Interest:** The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

The handling editor declared a shared affiliation with several of the authors YG, GA, and LZ at time of review.

Copyright © 2021 Gao, Ascoli and Zhao. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.