

PROBABILISTIC MODELING THE HIDDEN LAYERS OF DEEP NEURAL NETWORKS

Anonymous authors

Paper under double-blind review

ABSTRACT

In this paper, we demonstrate that the parameters of Deep Neural Networks (DNNs) cannot satisfy the i.i.d. prior assumption and activations being i.i.d. is not valid for all the hidden layers of DNNs. Hence, the Gaussian Process cannot correctly explain all the hidden layers of DNNs. Alternatively, we introduce a novel probabilistic representation for the hidden layers of DNNs in two aspects: (i) a hidden layer formulates a Gibbs distribution, in which neurons define the energy function, and (ii) the connection between two adjacent layers can be modeled by a product of experts model. Based on the probabilistic representation, we demonstrate that the entire architecture of DNNs can be explained as a Bayesian hierarchical model. Moreover, the proposed probabilistic representation indicates that DNNs have explicit regularizations defined by the hidden layers serving as prior distributions. Based on the Bayesian explanation for the regularization of DNNs, we propose a novel regularization approach to improve the generalization performance of DNNs. Simulation results validate the proposed theories.

1 INTRODUCTION

Recently, interpreting the hidden layers of Deep Neural Networks (DNNs) as the Gaussian Process (GP) has attracted a great deal of attention because it provides a novel probabilistic perspective to clarify the working mechanism of deep learning. Neal (1994) initially demonstrates the equivalence between GP and neural networks with a single fully connected layer in the limit of infinite neurons. Following this seminal work, Lee et al. (2018); Matthews et al. (2018) further extend the equivalence to neural networks with multiple hidden layers and Garriga-Alonso et al. (2018); Novak et al. (2018) establish the correspondence between the convolutional layers and GP.

It is important to note that the GP explanations rely on a fundamental probabilistic premise that the activations of a hidden layer are independent and identically distributed (i.i.d.). More specifically, based on the classical Central Limit Theorem (CLT), the prerequisite of an activation of a hidden layer approaching a Gaussian distribution is that all the activations of the previous layer are i.i.d.. Though the Lyapunov CLT could relax the probabilistic premise to be independence only, all the previous works do not thoroughly discuss the probabilistic premise except assuming an i.i.d. prior for all the parameters of DNNs. Due to its extreme importance but unclear status, it is necessary to thoroughly investigate the probabilistic premise for improving the interpretability of DNNs.

In this work, we demonstrate that the parameters of DNNs are correlated and depend on the training dataset, i.e., they cannot satisfy the i.i.d. prior. Moreover, we demonstrate that activations being i.i.d. cannot hold for all the hidden layers of DNNs. In the context of Bayesian probability, we theoretically derive the necessary conditions for the activations of a hidden layer being i.i.d. given the assumption that the activations of the previous layer are i.i.d.. Subsequently, we experimentally show that typical DNNs, such as the Multilayer Perceptron (MLP) and the Convolutional Neural Networks (CNNs), cannot satisfy the necessary conditions. As a result, activations being i.i.d. is not valid for all hidden layers. In other words, GP with the i.i.d. prior cannot correctly explain all the hidden layers of DNNs. In addition, some previous works show that GP is sensitive to the curse of dimensionality (Bengio et al., 2005; Hinton et al., 2012) and cannot clarify the hierarchical representation, an essential of deep learning (Matthews et al., 2018).

Alternatively, we propose a novel probabilistic representation for the hidden layers based on the Gibbs distribution (LeCun et al., 2006) and the Product of Experts (PoE) model (Hinton, 2002). The proposed probabilistic representation explains a hidden layer in two aspects: (i) a hidden layer can be formulated as a Gibbs distribution, in which neurons define the energy function, and (ii) the connection between two adjacent layers can be modeled by a PoE model. Specifically, the distribution of a single neuron or hidden unit (e.g., convolutional channel) can be expressed as a PoE, in which all the experts are defined by the neurons or hidden units in the previous hidden layer. Compared to the GP explanations, the proposed probabilistic representation provides a more specific explanation for the hidden layers of DNNs without any assumption.

Based on the probabilistic representation, the entire architecture of DNNs can be explained as a Bayesian Hierarchical Model (BHM) (Li & Perona, 2005). Precisely, the output layer formulates a likelihood distribution, the hidden layers recursively serve as prior distributions for their next layers, and the first hidden layer formulates a prior distribution of the training dataset. In contrast to previous works claiming that the backpropagation training performs implicit regularization for DNNs (Zhang et al., 2016; Neyshabur et al., 2017), the BHM explanation indicates that DNNs have explicit regularizations because prior distributions correspond to regularizations in the context of Bayesian regularization theory (Steck & Jaakkola, 2003). Based on the BHM explanations for the architecture of DNNs, we propose a novel regularization approach to improve the generalization performance of DNNs through pre-training the hidden layers corresponding to prior distributions.

2 BACKGROUND

Probabilistic modeling plays an extremely important role in machine learning. Most structured probabilistic models can be formulated as a Gibbs distribution, which describes the dependence of random variables \mathbf{x} by associating an energy to each dependence structure.

$$p(\mathbf{x}; \boldsymbol{\theta}) = \frac{1}{Z(\boldsymbol{\theta})} \exp[-E(\mathbf{x}; \boldsymbol{\theta})], \quad (1)$$

where $E(\mathbf{x}; \boldsymbol{\theta})$ is the energy function, $Z(\boldsymbol{\theta}) = \sum_{\mathbf{x}} \exp[-E(\mathbf{x}; \boldsymbol{\theta})]$ is the partition function, and $\boldsymbol{\theta}$ denote all the parameters. It is important to note that $E(\mathbf{x}; \boldsymbol{\theta})$ is the sufficient statistics for $p(\mathbf{x}; \boldsymbol{\theta})$, i.e., $p(\mathbf{x}; \boldsymbol{\theta})$ entirely depends on $E(\mathbf{x}; \boldsymbol{\theta})$ because $Z(\boldsymbol{\theta})$ only relies on $\boldsymbol{\theta}$ (Cover & Thomas, 2006). An important Gibbs distribution is the Markov Random Fields (MRFs) and the corresponding energy function is equivalent to the summation of multiple potential functions, i.e., $E(\mathbf{x}; \boldsymbol{\theta}) = -\sum_k f_k(\mathbf{x}; \boldsymbol{\theta}_k)$ (Geman & Geman, 1984; Wang et al., 2013; Goodfellow et al., 2016). Assuming a single potential function defines an expert, i.e., $F_k = \frac{1}{Z(\boldsymbol{\theta}_k)} \exp[-f_k(\mathbf{x}; \boldsymbol{\theta}_k)]$, we can view MRFs as a PoE model, i.e., $p(\mathbf{x}; \boldsymbol{\theta}) = \frac{1}{Z(\boldsymbol{\theta})} \prod_k F_k$, where $Z(\boldsymbol{\theta}) = \prod_k Z(\boldsymbol{\theta}_k)$ (Hinton, 2002).

Recently, numerous efforts have been devoted to explaining DNNs as a structured probabilistic models. These explanations include Boltzmann distribution (Salakhutdinov & Hinton, 2009; Mehta & Schwab, 2014), random field theory (Zheng et al., 2015), mixture model (Oord & Schrauwen, 2014; Tang et al., 2015; Patel et al., 2016) and Gaussian process (Williams, 1997; Borovykh, 2018).

3 MAIN RESULTS

3.1 PRELIMINARIES

We assume that \mathbf{X} and \mathbf{Y} are two random variables and $P_{\boldsymbol{\theta}}(\mathbf{X}, \mathbf{Y}) = P(\mathbf{Y}|\mathbf{X})p(\mathbf{X})$ is an unknown joint distribution between \mathbf{X} and \mathbf{Y} , where $P(\mathbf{X})$ describes the prior knowledge of \mathbf{X} , $P(\mathbf{Y}|\mathbf{X})$ describes the statistical connection between \mathbf{X} and \mathbf{Y} , and $\boldsymbol{\theta}$ indicate the parameters of $P_{\boldsymbol{\theta}}(\mathbf{X}, \mathbf{Y})$.

A dataset $\mathcal{D} = \{(\mathbf{x}^j, \mathbf{y}^j) | \mathbf{x}^j \in \mathbb{R}^M, \mathbf{y}^j \in \mathbb{R}^L\}_{j=1}^J$ is composed of i.i.d. samples generated from $P_{\boldsymbol{\theta}}(\mathbf{X}, \mathbf{Y})$. A neural network with I hidden layers is denoted as $\text{DNN} = \{\mathbf{x}; \mathbf{f}_1; \dots; \mathbf{f}_I; \mathbf{f}_Y\}$ and trained by \mathcal{D} , where $(\mathbf{x}, \mathbf{y}) \in \mathcal{D}$ are the input of the DNN and the corresponding training label, and \mathbf{f}_Y is an estimation of the true distribution $P(\mathbf{Y}|\mathbf{X})$. As a result, $\mathbf{x} \sim P(\mathbf{X})$ and $\mathbf{f}_Y \approx P(\mathbf{Y}|\mathbf{X})$. The meaning of \mathbf{f}_i is two-fold: (i) denoting all the neurons in the i th hidden layer, and (ii) representing the activations of the i th hidden layer. In addition, \mathbf{F}_i and \mathbf{F}_Y are the random variables corresponding to \mathbf{f}_i and \mathbf{f}_Y , respectively.




Figure 1: A MLP with two hidden layers, i.e., $DNN = \{\mathbf{x}; \mathbf{f}_1; \mathbf{f}_2; \mathbf{f}_Y\}$, in which \mathbf{f}_1 has N neurons, i.e., $\mathbf{f}_1 = \{f_{1n} = \sigma_1[g_{1n}(\mathbf{x})]\}_{n=1}^N$, where $g_{1n}(\mathbf{x}) = \sum_{m=1}^M \alpha_{mn} \cdot x_m + b_{1n}$ is the n th linear function with α_{mn} being the weight of the edge between x_m and f_{1n} and b_{1n} being the bias, and σ_1 is a non-linear activation function. In addition, $\mathbf{f}_2 = \{f_{2k} = \sigma_2[g_{2k}(\mathbf{f}_1)]\}_{k=1}^K$ where $g_{2k}(\mathbf{f}_1) = \sum_{n=1}^N \beta_{nk} \cdot f_{1n} + b_{2k}$ and $\mathbf{f}_Y = \{f_{Yl} = \sigma_3[g_{Yl}(\mathbf{f}_2)]\}_{l=1}^L$ where $g_{Yl}(\mathbf{f}_2) = \sum_{k=1}^K \gamma_{kl} \cdot f_{2k} + b_{Yl}$. They have the similar definitions as f_{1n} .

3.2 THE PARAMETERS OF DNNs CANNOT SATISFY THE I.I.D. ASSUMPTION

Above all, we examine the statistical properties of the parameters of DNNs on the benchmark MNIST and CIFAR-10 datasets and show the parameters of DNNs being correlated to each other. Since DNNs have so many parameters, it is very difficult to directly show the correlation of different parameters. Alternatively, we derive a necessary condition for the parameters of DNNs satisfying the i.i.d. prior and experimentally show that typical DNNs cannot satisfy the necessary condition. Therefore, the parameters of DNNs cannot satisfy the i.i.d. assumption.

If all the parameters of DNNs are i.i.d., the parameters of each neuron are identically distributed, e.g., $\{\alpha_{mn}\}_{m=1}^M \sim P(\alpha_n)$ in the MLP (Figure 1). Since all the parameters are i.i.d., different neurons should be uncorrelated, i.e., $\forall n \neq n', Corr(\alpha_n, \alpha_{n'}) = 0$. In particular, the i.i.d. assumption enables us to use the sample correlation $r(\alpha_n, \alpha_{n'})$ to estimate $Corr(\alpha_n, \alpha_{n'})$.

$$r(\alpha_n, \alpha_{n'}) = \frac{\sum_{m=1}^M (\alpha_{mn} - \bar{\alpha}_n)(\alpha_{mn'} - \bar{\alpha}_{n'})}{\sqrt{\sum_{m=1}^M (\alpha_{mn} - \bar{\alpha}_n)^2 \sum_{m=1}^M (\alpha_{mn'} - \bar{\alpha}_{n'})^2}} \approx Corr(\alpha_n, \alpha_{n'}) \quad (2)$$

where $\bar{\alpha}_n = \frac{1}{M} \sum_{m=1}^M \alpha_{mn}$ is the sample mean of $\{\alpha_{mn}\}_{m=1}^M$. Therefore, taking into account the estimation error, $|r(\alpha_n, \alpha_{n'})|$ close to zero is the necessary condition for all parameters being i.i.d..

First, we design a neural network with a single fully connected hidden layer, i.e., $NN = \{\mathbf{x}; \mathbf{f}; \mathbf{f}_Y\}$, to classify the MNIST dataset based on the seminal work (Neal, 1994). After training the NN well, we calculate the absolute sample correlations between the weights of different neurons in \mathbf{f} . The result shows that the weights do not satisfy the necessary condition in Appendix A.1. As a result, the parameters of the $NN = \{\mathbf{x}; \mathbf{f}; \mathbf{f}_Y\}$ cannot satisfy the i.i.d. assumption.

Second, we use a DNN (i.e., the above MLP) to classify the same dataset. The information about the network architecture and training methods is included in Appendix A.2. After training the MLP well, we calculate the absolute sample correlations, i.e., $|r(\alpha_n, \alpha_{n'})|$, $|r(\beta_k, \beta_{k'})|$, and $|r(\gamma_l, \gamma_{l'})|$, and derive three absolute correlation matrices, i.e., $A'_{N \times N}$, $B'_{K \times K}$, and $C'_{L \times L}$, shown in Figure 2. We find that most $|r(\alpha_n, \alpha_{n'})|$ are close to zero, but most $|r(\beta_k, \beta_{k'})|$ and $|r(\gamma_l, \gamma_{l'})|$ are non-zero. That indicates the necessary condition cannot hold for all the parameters of the MLP. In addition, we construct a CNN for classifying the CIFAR-10 dataset and examine the necessary condition for the CNN in Appendix A.3, which shows similar results as the MLP.

Third, we visualize all the parameters of the $NN = \{\mathbf{x}; \mathbf{f}; \mathbf{f}_Y\}$ for classifying a synthetic dataset in Appendix A.4. Since the synthetic dataset is much simpler than MNIST, we can directly show the correlation of different parameters and their dependences on the dataset by visualizing them. Overall, these simulations demonstrate that the parameters of DNNs cannot be i.i.d..




Figure 2: Three absolute correlation matrices for the weights in \mathbf{f}_1 , \mathbf{f}_2 , and \mathbf{f}_Y . In addition, 0.08, 0.26, 0.31 is the average of the absolute correlation coefficients for all the weights in \mathbf{f}_1 , \mathbf{f}_2 , and \mathbf{f}_Y , respectively.

3.3 ACTIVATIONS BEING I.I.D. IS NOT VALID FOR ALL THE HIDDEN LAYERS OF DNNs

Although we demonstrate that the parameters of DNNs cannot satisfy the i.i.d. assumption, we are still unclear whether the activations of a hidden layer are i.i.d.. In this section, we demonstrate that activations being i.i.d. is not valid for all the hidden layers of DNNs.

In the context of Bayesian probability, all the weights and biases of the MLP (Figure 1) have prior distributions, i.e., $\alpha_{mn} \sim P(A_{mn})$, $\beta_{nk} \sim P(W_{nk})$, $b_{1n} \sim P(B_{1n})$, and $b_{2k} \sim P(B_{2k})$, thus we regard $G_{1n} = \sum_{m=1}^M A_{mn}X_m + B_{1n}$ as the random variable of $g_{1n}(\mathbf{x}) = \sum_{m=1}^M \alpha_{mn} \cdot x_m + b_{1n}$, and $G_{2k} = \sum_{n=1}^N W_{nk}F_{1n} + B_{2k}$ as the random variable of $g_{2k}(\mathbf{f}_1) = \sum_{n=1}^N \beta_{nk} \cdot f_{1n} + b_{2k}$. Hence, the random variables of two arbitrary activations in \mathbf{f}_1 and \mathbf{f}_2 , e.g., f_{1n} and f_{2k} , can be expressed as $F_{1n} = \sigma_1(G_{1n})$ and $F_{2k} = \sigma_2(G_{2k})$, respectively.

Given two random variables A and B , the necessary conditions for them being i.i.d. are that they are uncorrelated (i.e., $Cov(A, B) = 0$) and have the same expectation (i.e., $E(A) = E(B)$). Similarly, assuming $\{F_{1n}\}_{n=1}^N$ are i.i.d. and $\{f_{1n}\}_{n=1}^N \sim P(F_1)$, we derive the following necessary conditions for the activations of the second hidden layer, i.e., $\{F_{2k}\}_{k=1}^K$, being i.i.d..

$$\forall (k, k') \in S_1 = \{(k, k') \in \mathbb{Z}^2 | k \neq k', 1 \leq k \leq K, 1 \leq k' \leq K\}, \text{ we must have}$$

$$Cov(G_{2k}, G_{2k'}) = Cov(W_k, W_{k'})E^2(F_1) + Cov(W_k, B_{2k'})E(F_1) + Cov(W_{k'}, B_{2k})E(F_1) + Cov(B_{2k}, B_{2k'}) = 0 \quad (3)$$

$$E(F_1)E(W_k) + E(B_{2k}) = E(F_1)E(W_{k'}) + E(B_{2k'}) \quad (4)$$

$$\sigma_2(\cdot) \text{ is strictly increasing, differentiable, and invertible.} \quad (5)$$

where $W_k = \sum_{n=1}^N W_{nk}$. Equation 3 and 4 are the necessary conditions for $\{G_{2k}\}_{k=1}^K$ being i.i.d.. Equation 5 constrains the activation function $\sigma_2(\cdot)$ to guarantee that if $\{G_{2k}\}_{k=1}^K$ are i.i.d. then $\{F_{2k}\}_{k=1}^K$ are i.i.d. as well. The necessary conditions are valid for arbitrary fully connected layers as long as properly changing the subscripts. The detailed derivations are included in Appendix B and Appendix C. Moreover, we prove that the necessary conditions also hold for convolutional layers based on the connection between fully connected layers and convolutional layers in Appendix E.

It is noteworthy that the i.i.d. assumption for the parameters of DNNs indeed satisfy the necessary conditions for activations being i.i.d.. Nevertheless, since the i.i.d. prior is not an appropriate prior for the parameters of DNNs, it cannot guarantee activations of a hidden layer being i.i.d..

Alternatively, the theoretical necessary conditions indicate an applicable approach to examine whether activations of a hidden layer being i.i.d. given the assumption that the activations of the previous layer are i.i.d.. More specifically, since $\{F_{1n}\}_{n=1}^N$ are i.i.d., we can use the sample mean $\bar{\mathbf{f}}_1$ to estimate $E(F_1)$. Because the samples of the training dataset \mathcal{D} are also assumed to be i.i.d., we can train the MLP multiple times to draw multiple independent observations of W_{nk} and B_{2k} , thus we can use the sample covariance to estimate $Cov(W_k, W_{k'})$, $Cov(W_k, B_{2k'})$, and etc.

Table 1: The statistical measures for examining the necessary conditions

Layer	Activation expectation	Covariance	Weight expectation
\mathbf{f}_1	$E(X)$	$Cov(A_n, A_{n'})$	$E(A_n)$
\mathbf{f}_2	$E(F_1)$	$Cov(W_k, W_{k'})$	$E(W_k)$
\mathbf{f}_Y	$E(F_2)$	$Cov(C_l, C_{l'})$	$E(C_l)$

$$\{x_m\}_{m=1}^M \sim P(X), \{f_{1n}\}_{n=1}^N \sim P(F_1), \{f_{2k}\}_{k=1}^K \sim P(F_2)$$

$$A_n = \sum_{m=1}^M A_{mn}, \text{ where } \alpha_{mn} \sim P(A_{mn})$$

$$W_k = \sum_{n=1}^N W_{nk}, \text{ where } \beta_{nk} \sim P(W_{nk})$$

$$C_l = \sum_{k=1}^K C_{kl}, \text{ where } \gamma_{kl} \sim P(C_{kl})$$

Based on the empirical approach, we estimate the statistical measures (Table 1) to examine whether the activations of the three layers (i.e., \mathbf{f}_1 , \mathbf{f}_2 , and \mathbf{f}_Y) being i.i.d. in the MLP for classifying the MNIST dataset. Without loss of generality, we restrict all the layers of the MLP from using bias to decrease computation complexity, thus the necessary conditions for \mathbf{f}_2 can be simplified as

$$\forall (k, k') \in S_1 = \{(k, k') \in \mathbb{Z}^2 | k \neq k', 1 \leq k \leq K, 1 \leq k' \leq K\}, \text{ we must have} \quad (6)$$

$$Cov(G_{2k}, G_{2k'}) = Cov(W_k, W_{k'})E^2(F_1) = 0$$

$$E(F_1)E(W_k) = E(F_1)E(W_{k'}) \quad (7)$$

First, Figure 3 shows the correlation matrices, i.e., $A_{N \times N}$, $W_{K \times K}$, and $C_{L \times L}$, to check if their corresponding covariances, i.e., $Cov(A_n, A_{n'})$, $Cov(W_k, W_{k'})$, and $Cov(C_k, C_{k'})$, are close to zero. It shows that many correlations are far from zero, so their corresponding covariances are not close to zero either. Second, we use the sample mean, e.g., \bar{x} , to estimate the expectation, e.g., $E(X)$. Taking into account the estimation error, we can regard $E(X) = 0$ based on $\bar{x} = 0.131$. However, $\bar{\mathbf{f}}_1 = 0.436$ and $\bar{\mathbf{f}}_2 = 0.498$ imply $E(F_1) \neq 0$ and $E(F_2) \neq 0$ even considering the estimation error. As a result, the activations of the last two hidden layer, i.e., \mathbf{f}_2 and \mathbf{f}_Y , cannot be independent because they cannot satisfy the first necessary condition, i.e., Equation 6. The detailed information about the estimation of the statistical measures and the architecture of the MLP are presented in Appendix F.1. In addition, we show that the necessary condition for activations being identically distributed also cannot be satisfied by the MLP in Appendix F.1. Furthermore, we demonstrate that the necessary conditions for activations being i.i.d. also cannot be satisfied in the CNN for classifying the CIFAR-10 dataset in Appendix F.2

Overall, we conclude that activations being i.i.d. is not valid for all the hidden layers of DNNs based on the theoretical necessary conditions and the experimental simulations on typical DNNs. In other words, we cannot use the CLT to establish the equivalence between GP and all the hidden layers of DNNs, thus GP with the i.i.d. prior could not correctly clarify all the hidden layers of DNNs. That necessitates a new probabilistic representation for explaining the hidden layers of DNNs.




Figure 3: Three absolute correlation matrices for showing whether $Cov(A_n, A_{n'})$, $Cov(W_k, W_{k'})$, and $Cov(C_k, C_{k'})$ are close to zero.

3.4 THE DISTRIBUTION OF A HIDDEN LAYER IS A GIBBS DISTRIBUTION

This section proves that the distribution of a hidden layer can be formulated as a Gibbs distribution and the statistical connection between two adjacent layers can be modeled as a PoE model.

Assuming the output layer in the MLP (Figure 1) is the softmax, its distribution can be expressed as

$$P(\mathbf{F}_Y) = \{P(f_{yl}) = \frac{1}{Z_Y} \exp(f_{yl})\}_{l=1}^L \quad (8)$$

where $Z_Y = \sum_{l=1}^L \exp(f_{yl})$ is the partition function. Since $f_{yl} = \sum_{k=1}^K \gamma_{kl} \cdot f_{2k} + b_{yl}$, we have

$$P(\mathbf{F}_Y) = \{P(f_{yl}) = \frac{1}{Z_Y} \exp(\sum_{k=1}^K \gamma_{kl} \cdot f_{2k} + b_{yl})\}_{l=1}^L \quad (9)$$

Based on the properties of the exponential function, i.e., $\exp(a + b) = \exp(a) \cdot \exp(b)$ and $\exp(a \cdot b) = [\exp(b)]^a$, we can reformulate $P(\mathbf{F}_Y)$ as

$$P(\mathbf{F}_Y) = \{P(f_{yl}) = \frac{1}{Z'_Y} \prod_{k=1}^K [\exp(f_{2k})]^{\gamma_{kl}}\}_{l=1}^L \quad (10)$$

where $Z'_Y = Z_Y / \exp(b_{yl})$. Since $\{\exp(f_{2k})\}_{k=1}^K$ are scalar, we can introduce a new partition function $Z_{F_2} = \sum_{k=1}^K \exp(f_{2k})$ such that $\{\frac{1}{Z_{F_2}} \exp(f_{2k})\}_{k=1}^K$ become a probability measure. As a result, we can further reformulate $P(\mathbf{F}_Y) = \{P(f_{yl})\}_{l=1}^L$ as a Product of Expert (PoE) model as

$$P(\mathbf{F}_Y) = \{P(f_{yl}) = \frac{1}{Z''_Y} \prod_{k=1}^K [\frac{1}{Z_{F_2}} \exp(f_{2k})]^{\gamma_{kl}}\}_{l=1}^L \quad (11)$$

where $Z''_Y = Z_Y / [\exp(b_{yl}) \cdot \prod_{k=1}^K [Z_{F_2}]^{\gamma_{kl}}]$ and each expert is defined as $\frac{1}{Z_{F_2}} \exp(f_{2k})$.

It is noteworthy that all the experts $\{\frac{1}{Z_{F_2}} \exp(f_{2k})\}_{k=1}^K$ form a probability measure and establish an exact one-to-one correspondence to all the neurons in the second hidden layer \mathbf{f}_2 , i.e., $\{f_{2l}\}_{k=1}^K$. Therefore, the distribution of \mathbf{f}_2 can be expressed as

$$P(\mathbf{F}_2) = \{P(f_{2k}) = \frac{1}{Z_{F_2}} \exp(f_{2k})\}_{k=1}^K \quad (12)$$

Based on the definition of Gibbs distribution, Equation 8 and 12 show that \mathbf{f}_Y and \mathbf{f}_2 formulate two multivariate Gibbs distributions and their energy functions are equivalent to the negative of the output nodes $\{f_{yl}\}_{l=1}^L$ and the neurons $\{f_{2k}\}_{k=1}^K$, respectively. In addition, Equation 11 implies that the connection between $P(f_{yl})$ and $P(f_{2k})$ can be expressed as the PoE model, namely that $P(f_{yl})$ is equal to the product of all the experts $\{P(f_{2k})\}_{k=1}^K$ with their respective weights $\{\gamma_{kl}\}_{k=1}^K$.

Compared to the GP explanations assuming all the activations $\{f_{2k}\}_{k=1}^K$ are i.i.d. and deriving $P(f_{yl})$ being a Gaussian distribution in the limit of infinite neurons, i.e., $\{f_{2k}\}_{k=1}^K \rightarrow \infty$, based on CLT, we formulate $P(f_{yl})$ as a PoE model, in which all the experts $\{\frac{1}{Z_{F_2}} \exp(f_{2k})\}_{k=1}^K$ are defined by all the neurons $\{f_{2k}\}_{k=1}^K$. Overall, the above probabilistic representation provides a more specific explanation for the hidden layer without any assumption.

Moreover, we prove that the probabilistic representation is valid for other hidden layers in the MLP. Since $\{f_{2k} = \sigma_2(\sum_{n=1}^N \beta_{nk} \cdot f_{1n} + b_{2k})\}_{k=1}^K$, $P(\mathbf{F}_2)$ can be expressed as

$$P(\mathbf{F}_2) = \{P(f_{2k}) = \frac{1}{Z_{F_2}} \exp[\sigma_2(\sum_{n=1}^N \beta_{nk} \cdot f_{1n} + b_{2k})]\}_{k=1}^K \quad (13)$$

Based on the equivalence between the gradient descent algorithm and the first order approximation (Battiti, 1992), we prove that $\sigma_2(\sum_{n=1}^N \beta_{nk} \cdot f_{1n} + b_{2k})$ can be approximated as

$$\sigma_2(\sum_{n=1}^N \beta_{nk} \cdot f_{1n} + b_{2k}) \approx C_{21} \cdot [\sum_{n=1}^N \beta_{nk} \cdot f_{1n} + b_{2k}] + C_{22} \quad (14)$$

where C_{21} and C_{22} only depend on the activations $\{f_{1n}\}_{n=1}^N$ in the previous training iteration, thus they can be regarded as constants and absorbed by β_{nk} and b_{2k} . The proof for the approximation is included in Appendix G. Therefore, $P(\mathbf{F}_2)$ also can be modeled as a PoE model.

$$P(\mathbf{F}_2) = \{P(F_{2k}) \approx \frac{1}{Z''_{\mathbf{F}_2}} \prod_{n=1}^N [\frac{1}{Z_{\mathbf{F}_1}} \exp(f_{1n})]^{\beta_{nk}}\}_{k=1}^K \quad (15)$$

where $Z''_{\mathbf{F}_2} = Z_{\mathbf{F}_2} / [\exp(b_{2k}) \cdot \prod_{n=1}^N [Z_{\mathbf{F}_1}]^{\beta_{nk}}]$ and $Z_{\mathbf{F}_1} = \sum_{n=1}^N \exp(f_{1n})$.

Finally, the distribution of the first hidden layer \mathbf{f}_1 is formulated as

$$P(\mathbf{F}_1) = \{P(F_{1n}) = \frac{1}{Z_{\mathbf{F}_1}} \exp(f_{1n})\}_{n=1}^K \quad (16)$$

Based on the same derivation as $P(\mathbf{F}_2)$, $P(\mathbf{F}_1)$ also can be modeled as a PoE model as follows.

$$P(\mathbf{F}_1) = \{P(F_{1n}) \approx \frac{1}{Z''_{\mathbf{F}_1}} \prod_{m=1}^M [\frac{1}{Z_{\mathbf{X}}} \exp(x_m)]^{\alpha_{mn}}\}_{n=1}^N \quad (17)$$

where $Z''_{\mathbf{F}_1} = Z_{\mathbf{F}_1} / [\exp(b_{1n}) \cdot \prod_{m=1}^M [Z_{\mathbf{X}}]^{\alpha_{mn}}]$ and $Z_{\mathbf{X}} = \sum_{m=1}^M \exp(x_m)$.

In summary, we introduce a novel probabilistic representation for the hidden layers in two aspects. First, a fully connected layer formulate a multivariate Gibbs distribution (Equation 8, Equation 12, and Equation 16) to describe the probability of the occurrence of multiple features formulated by their respective neurons and their energy functions are equivalent to the negative of their respective neurons. Second, the statistical connection between two adjacent fully connected layers can be expressed by a PoE model (Equation 11, Equation 15, and Equation 17).

Furthermore, we prove that a convolutional layer formulates a specific Gibbs distribution, i.e., the MRF model, in Appendix H. Specifically, a convolutional layer \mathbf{f}_i with K convolutional channels, i.e., $\mathbf{f}_i = \{\mathbf{f}_i^k\}_{k=1}^K$, in a DNN $= \{\mathbf{x}; \mathbf{f}_1; \dots; \mathbf{f}_I; \mathbf{f}_Y\}$ can be formulated as a MRF model.

$$P(\mathbf{F}_i) = \frac{1}{Z_{\mathbf{F}_i}} \prod_{k=1}^K \exp(\mathbf{f}_i^k) = \frac{1}{Z_{\mathbf{F}_i}} \exp(\sum_{k=1}^K \mathbf{f}_i^k) \quad (18)$$

where $\mathbf{f}_i^k = \sum_{n=1}^N \sigma_i(\mathbf{S}_i^{(k,n)} \circ \mathbf{f}_{i-1}^n) + b_i^k \cdot \mathbf{1}$ is the k th convolutional channel derived from the previous layer with N channels, i.e., $\mathbf{f}_{i-1} = \{\mathbf{f}_{i-1}^n\}_{n=1}^N$, $\mathbf{S}_i^{(k,n)}$ is the convolutional filter connecting \mathbf{f}_i^k and \mathbf{f}_{i-1}^n , b_i^k is the bias, and $\sigma_i(\cdot)$ is the activation function. In particular, its energy function is equal to the negative summation of all the convolutional channels, i.e., $E_{\mathbf{F}_i} = -\sum_{k=1}^K \mathbf{f}_i^k$

4 SIMULATIONS

This section uses two synthetic datasets to demonstrate the proposed probabilistic representation: (i) the distribution of a fully connected layer is a multivariate discrete Gibbs distribution and (ii) the distribution a convolutional layer is a MRF model.

4.1 THE DISTRIBUTION OF A FULLY CONNECTED LAYER IS A MULTIVARIATE DISCRETE GIBBS DISTRIBUTION

A fully connected layer $\mathbf{f} = \{f_n\}_{n=1}^N$ with N neurons formulates a multivariate Gibbs distribution to describe the probability of the N features defined the neurons occurring in the input. Their energy functions are equal to the negative of their respective neurons, i.e., $E_{\mathbf{f}_n} = -f_n(\mathbf{x})$.

$$P(\mathbf{F}) = \{P(f_n) = \frac{1}{Z_{\mathbf{F}}} \exp[f_n(\mathbf{x})]\}_{n=1}^N \quad (19)$$

where $f_n(\mathbf{x}) = \sigma[g_n(\mathbf{x})]$, $g_n(\mathbf{x}) = \sum_{m=1}^M \alpha_{mn} \cdot x_m + b_n$ is a linear filter with α_{mn} being the weights and b_n being the bias, $\sigma(\cdot)$ is the activation function, and $Z_{\mathbf{F}} = \sum_{n=1}^N \exp[f_n(\mathbf{x})]$ is the partition function.




Figure 4: The top-left figure shows a 32×32 synthetic image as the input, which is sampled from the Gaussian distribution $\mathcal{N}(0, 1)$ and sorted in the primary diagonal direction by the descending order. The remaining figures visualize all the weights $\{\alpha_{mn}\}_{m=1}^{1024}$ of the seven neurons. In addition, the subscript in each figure, i.e., $(g_n(\mathbf{x}), f_n(\mathbf{x}), P(f_n))$, denotes the linear output $g_n(\mathbf{x})$, activation $f_n(\mathbf{x})$, and Gibbs probability $P(f_n)$.

Since benchmark datasets consist of very complex features and typical DNNs have too complicated architecture, it is very hard to directly derive the statistical connection between the training dataset and the learned parameters of DNNs. Alternatively, we choose the neural network for classifying the synthetic dataset in Appendix A.4 to demonstrate the probabilistic representation, because the neural network only has a single fully connected hidden layer with seven neurons, i.e., $N = 7$, and the synthetic dataset only has four features. After the network is trained well, we choose a synthetic image as the input and show all the information of the learned hidden layer in Figure 4.

Since the synthetic dataset contains four different features in the primary and secondary diagonal directions, the neurons are shown to learn seven filters to distinguish these features. Also note that the input image is sorted in the primary diagonal direction, i.e., it has a huge difference between the two sides of the secondary diagonal (the red curve), thus Neuron1, Neuron 2 and Neuron 4 derive large positive output (i.e. $g_1(\mathbf{x}) = 122.09$, $g_2(\mathbf{x}) = 113.70$, and $g_4(\mathbf{x}) = 140.20$), but Neuron6 and Neuron 7 derive large negative output (i.e. $g_6(\mathbf{x}) = -134.08$, $g_7(\mathbf{x}) = -133.36$). After applying the sigmoid function, we have $f_1(\mathbf{x}) = f_2(\mathbf{x}) = f_4(\mathbf{x}) = 1$ and $f_6(\mathbf{x}) = f_7(\mathbf{x}) = 0$. Based on the definition of Gibbs distribution (Equation 19), the Gibbs probability of different neurons are equal to $P(f_1) = P(f_2) = P(f_4) = 0.20$ and $P(f_6) = P(f_7) = 0.07$, which indicates that the probability of the occurrence of the features defined by Neuron1, Neuron2 and Neuron4 is high, but the probability corresponding to Neuron6 and Neuron 7 are very low in the synthetic image.

4.2 THE DISTRIBUTION OF A CONVOLUTIONAL LAYER IS A MRF MODEL

Since the distributions of benchmark datasets are unknown, it is impossible to use them to verify the probabilistic explanation for convolutional layers. Alternatively, we generate a synthetic dataset obeying the Gaussian distribution $\mathcal{N}(0, 1024)$ based on the NIST dataset of handwritten digits. It consists of 20,000 32×32 grayscale images in 10 classes (digits from 0 to 9), and each class has 1,000 training images and 1,000 testing images. Figure 5 shows five synthetic images and their perspective histograms. The method for generating the synthetic dataset is included in Appendix I. We design a CNN = $\{\mathbf{x}; \mathbf{f}_1; \mathbf{f}_2; \mathbf{f}_3; \mathbf{f}_4; \mathbf{f}_Y\}$ for classifying the synthetic dataset. The CNN has two convolutional layers with ReLU (i.e., \mathbf{f}_1 and \mathbf{f}_3), two max-pooling layers (i.e., \mathbf{f}_2 and \mathbf{f}_4) and a softmax output layer \mathbf{f}_Y . The detailed information about the CNN is included in Appendix I.

Based on the MRF explanation (Equation 18), the distribution of \mathbf{f}_1 can be formulated as

$$P(\mathbf{F}_1) = \frac{1}{Z_{\mathbf{F}_1}} \prod_{k=1}^{20} \exp(\mathbf{f}_1^k) = \frac{1}{Z_{\mathbf{F}_1}} \exp\left(\sum_{k=1}^{20} \mathbf{f}_1^k\right) \quad (20)$$

In particular, the energy function of $P(\mathbf{F}_1)$ is equivalent to $E_{\mathbf{F}_1} = -\sum_{k=1}^{20} \mathbf{f}_1^k$.

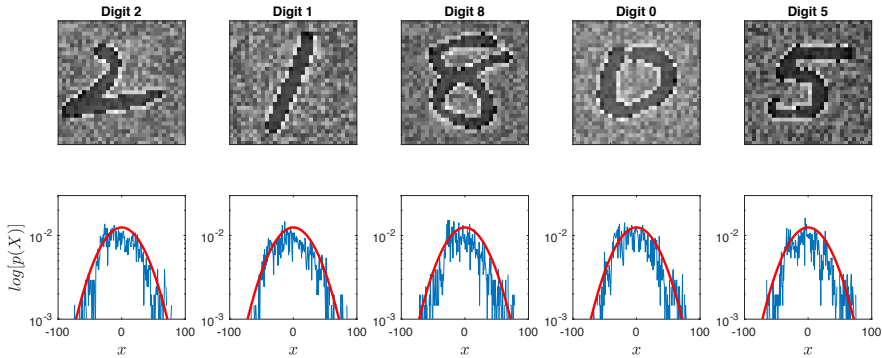


Figure 5: The first row shows five synthetic images of handwritten digits, the second row visualizes their respective histograms, and the red curve indicates the Gaussian distribution $\mathcal{N}(0, 1024)$.

In order to model a high dimensional dataset, e.g., the 32×32 synthetic image, we typically need a multivariate distribution with the same dimension, e.g., the dimension of the covariance function of a Gaussian process should be 1024×1024 . However, most structured probabilistic graphical models, especially the MRF model, have two fundamental assumptions: stationary and Markov (Lyu & Simoncelli, 2007). The former means that the distribution is independent on the location, i.e., the distribution of pixels at different locations are identically distributed. The later indicates that the distribution of a single pixel is independent on the other pixels given its neighbors. The two assumptions allow us to use the histogram of samples to simulate the true distribution.




Figure 6: The red curve indicates the truly prior distribution $P(\mathbf{X}) = \mathcal{N}(0, 1024)$. The blue curves are different histograms. (A) the synthetic image \mathbf{x} is the input. (B) the histogram of \mathbf{x} , i.e., $P(\mathbf{x})$, and $\text{KL}[P(\mathbf{X})||P(\mathbf{x})] = 0.62$. (C) the histogram of \mathbf{E}_{F_1} for estimating $P(F_1)$ and $\text{KL}[P(\mathbf{X})||P(F_1)] = 0.83$.

After the CNN is well trained, we randomly choose a testing image \mathbf{x} as the input to derive $P(F_1)$. Since $\mathbf{x} \sim P(\mathbf{X})$, we can use the histogram of the testing image, i.e., $P(\mathbf{x})$, to estimate $P(\mathbf{X})$. In addition, \mathbf{E}_{F_1} is a sufficient statistics for $P(F_1)$ that allows us to use the histogram of \mathbf{E}_{F_1} to estimate $P(F_1)$. As a result, we can verify the probabilistic representation by calculating the distance between $P(F_1)$ and $P(\mathbf{X})$, i.e., $\text{KL}[P(\mathbf{X})||P(F_1)]$. We visualize their distributions in Figure 6, which shows that $P(F_1)$ is very close to $P(\mathbf{X})$, i.e., $\text{KL}(P(\mathbf{X})||q(F_1)) = 0.83$. Therefore, the distribution of a convolutional layer can be modeled as a MRF model.

In addition, this experiment provides a viewpoint to show the activations of different convolutional layers not being i.i.d.. Given another synthetic image, we visualize the output of five convolutional channels in \mathbf{f}_1 in Figure 7. It shows that the activations of different channels at the same position are correlated, i.e., they cannot be independent. In addition, we estimate the distribution of each convolutional channel based on the histograms of their respective energy functions. It shows that the distributions of different channels are different, i.e., they cannot be identically distributed.

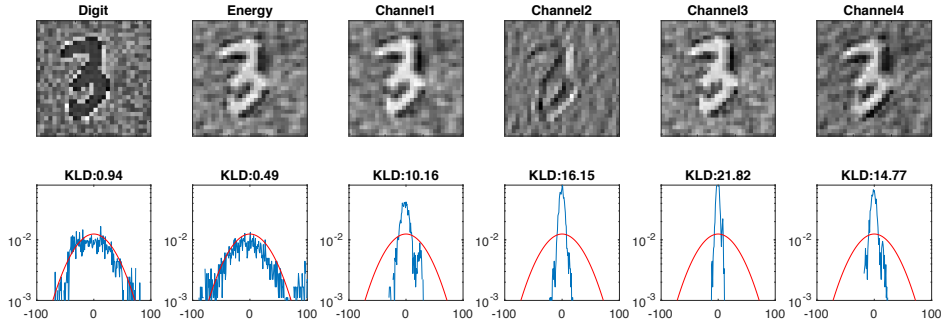


Figure 7: The first row shows a synthetic image, the energy functions of $P(\mathbf{F}_1)$ and five different experts, i.e., $\{P(\mathbf{F}_1^n) = \frac{1}{Z_{\mathbf{F}_1^n}} \exp(\mathbf{f}_1^n)\}_{n=1}^5$. The second row shows their respective distributions, and the KLDs denote their distances to the true distribution, i.e., $\text{KLD}[P(\mathbf{X})||P(\mathbf{F}_1^n)]$.

5 DISCUSSION

5.1 THE ENTIRE ARCHITECTURE OF DNNs CAN BE EXPLAINED AS A BHM

Since the output of the hidden layer \mathbf{f}_i in a DNN $= \{\mathbf{x}; \mathbf{f}_1; \dots; \mathbf{f}_I; \mathbf{f}_Y\}$ is commonly the input of the next layer \mathbf{f}_{i+1} , the DNN forms a Markov chain as $\mathbf{F}_1 \rightarrow \dots \rightarrow \mathbf{F}_I \rightarrow \mathbf{F}_Y$. Hence, the distribution of a hidden layer in DNNs should be viewed as a conditional distribution. As a result, the entire architecture of DNNs formulates a joint distribution as

$$P(\mathbf{F}_Y; \mathbf{F}_I; \dots; \mathbf{F}_1 | \mathbf{X}) = P(\mathbf{F}_Y | \mathbf{F}_I) \cdot \dots \cdot P(\mathbf{F}_{i+1} | \mathbf{F}_i) \cdot \dots \cdot P(\mathbf{F}_1 | \mathbf{X}) \quad (21)$$

Since $P(\mathbf{F}_Y | \mathbf{F}_I)$ is an estimation of the true distribution $P(\mathbf{Y} | \mathbf{X})$, it can be regarded as a likelihood distribution, and the simulations show that $P(\mathbf{F}_1 | \mathbf{X})$ can be regarded as a prior distribution of the training dataset. The Markov chain indicates that other hidden layers recursively serve as prior distributions for their respective next layers, i.e., $P(\mathbf{F}_{i+1} | \mathbf{F}_i)$. Therefore, the entire architecture of DNNs can be explained as a BHM, which provides a novel probabilistic explanation for the hierarchy property of DNNs, which cannot be clarified by GP (Matthews et al., 2018)

5.2 A NOVEL REGULARIZATION ALGORITHM FOR DNNs

In contrast to recent works claim that DNNs perform implicit regularization by the backpropagation learning (Zhang et al., 2016; Neyshabur et al., 2017), the BHM explanation indicates that DNNs have explicit regularizations represented by the hidden layers serving as prior distributions based on the Bayesian regularization theory (Steck & Jaakkola, 2003). This novel explanation for DNNs indicates a new regularization approach to improve the generalization performance of DNNs through pre-training the hidden layers corresponding to prior distributions. In particular, we introduce a novel regularization algorithm for CNNs and demonstrate its superiority over traditional regularizations, e.g., dropout (Srivastava et al., 2014), based on the CIFAR-10 classification task in Appendix J.

6 CONCLUSION

In this paper, we demonstrate that the parameters of DNN cannot satisfy the i.i.d. assumption and the activations being i.i.d. cannot be valid for all the hidden layers of DNNs. Hence, GP cannot correctly explain all the hidden layer of DNNs. Alternatively, we propose a novel probabilistic representation for DNNs in three aspects: (i) a hidden layer formulates a Gibbs distribution, in which neurons define the energy function, (ii) the connection between two adjacent layers can be modeled by a PoE model, and (iii) the entire architecture of DNNs can be explained as a BHM.

There are two general directions for future research. The theoretical direction is to examine the probabilistic representation in other popular DNNs, e.g., the residual networks, which can further verify the proposed probabilistic representation. The practical direction is to validate the proposed regularization in more complex DNNs and datasets.

REFERENCES

- Roberto Battiti. First and second order methods for learning: Between steepest descent and newton's method. *Neural Computation*, 4:141–166, 1992.
- Yoshua Bengio, Olivier Delalleau, and Nicolas Le Roux. The curse of dimensionality for local kernel machines. *Technical Report 1258*, 2005.
- Anastasia Borovykh. A gaussian process perspective on convolutional neural networks. In *arXiv:1810.10798*, 2018.
- M. A. Carreira-Perpinan and Geoffrey E. Hinton. On contrastive divergence learning. *Artificial Intelligence and Statistics*, pp. 33–41, 2005.
- Thomas Cover and Joy Thomas. *Elements of Information Theory*. Wiley-Interscience, Hoboken, New Jersey, 2006.
- Adri Garriga-Alonso, Carl Edward Rasmussen, and Laurence Aitchison. Deep convolutional networks as shallow gaussian processes. In *ICLR*, 2018.
- S. Geman and D. Geman. Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 721–741, June 1984.
- Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- K He, X Zhang, S Ren, and J Sun. Deep residual learning for image recognition. In *CVPR*, pp. 770–778, 2016.
- G. Hinton, L. Deng, D. Yu, G. Dahl, A. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. Sainath, and B. Kingsbury. Deep neural networks for acoustic modeling in speech recognition. *IEEE Signal Processing Magazine*, 2012.
- Geoffrey E. Hinton. Training products of experts by minimizing contrastive divergence. *Neural Computation*, 14:1771–1800, 2002.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. *Technique Report*, 2009.
- Y. LeCun, L. Bottou, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 11:2278–2324, 1998.
- Yann LeCun, Sumit Chopra, Raia Hadsell, Marc’Aurelio Ranzato, and Fu Jie Huang. *A tutorial on energy-based learning*. MIT Press, 2006.
- Jaehoon Lee, Yasaman Bahri, Roman Novak, Samuel S. Schoenholz, Jeffrey Pennington, and Jascha Sohl-Dickstein. Deep neural networks as gaussian processes. In *ICLR*, 2018.
- Feifei Li and Pietro Perona. A bayesian hierarchical model for learning natural scene categories. In *CVPR*, 2005.
- Siwei Lyu and Piero P. Simoncelli. Statistical modeling of images with fields of gaussian scale mixtures. In *NeurIPS*, 2007.
- Alexander Matthews, Mark Rowland, Jiri Hron, Richard E. Turner, and Zoubin Ghahramani. Gaussian process behaviour in wide deep neural networks. In *ICLR*, 2018.
- Pankaj Mehta and David J. Schwab. An exact mapping between the variational renormalization group and deep learning. *arXiv preprint arXiv:1410.3831*, 2014.
- Radford M. Neal. Priors for infinite networks. *Technical report no. crg-tr-94-1*, 1994.
- Behnam Neyshabur, Srinadh Bhojanapalli, David McAllester, and Nathan Srebro. Exploring generalization in deep learning. In *NeurIPS*, 2017.

- Roman Novak, Lechao Xiao, Jaehoon Lee, Yasaman Bahri, Greg Yang, Jiri Hron, Daniel A. Abolafia, Jeffrey Pennington, and Jascha Sohl-Dickstein. Bayesian deep convolutional networks with many channels are gaussian processes. In *ICLR*, 2018.
- Aaron Oord and B Schrauwen. Factoring variations in natural images with deep gaussian mixture models. In *NeurIPS*, 2014.
- Ankit Patel, Minh Nguyen, and Richard Baraniuk. A probabilistic framework for deep learning. In *NeurIPS*, 2016.
- Stefan Roth and Michael J. Black. Fields of experts: A framework for learning image priors. *IEEE Conf. on CVPR*, pp. 860–867, July 2005.
- Stefan Roth and Michael J. Black. Fields of experts. *International Journal of Computer Vision*, 82: 205–229, 2008.
- David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. Learning representations by back-propagating errors. *Nature*, 323:533–536, October 1986.
- Ruslan Salakhutdinov and Geoffrey Hinton. Deep boltzmann machines. In *AISTATS 2009*, pp. 448–455, 2009.
- U. Schmidt, K. Schelten, and S. Roth. Bayesian deblurring with integrated noise estimation. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pp. 2625–2632, 2011.
- Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15:1929–1958, 2014.
- Harald Steck and Tommi S. Jaakkola. On the dirichlet prior and bayesian regularization. In *NeurIPS*, 2003.
- Yichuan Tang, Ruslan Salakhutdinov, and Geoffrey Hinton. Deep mixtures of factor analysers. *arXiv preprint arXiv:1206.4635*, 2015.
- C. Wang, N. Komodakis, and N. Paragios. Markov random field modeling inference and learning in computer vision and image understanding: A survey. *Computer Vision and Image Understanding*, 117:1610–1627, 2013.
- Christopher KI Williams. Computing with infinite networks. In *NeurIPS*, pp. 295–301, 1997.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *ICLR*, 2016.
- Shuai Zheng, Sadeep Jayasumana, Bernardino Romera-Paredes, Vibhav Vineet, Zhizhong Su, Dalong Du, Chang Huang, and Philip Torr. Conditional random fields as recurrent neural networks. In *International Conference on Computer Vision (ICCV)*, pp. 1529–1537, 2015.

Table 2: The architecture of the NN = { \mathbf{x} ; \mathbf{f} ; \mathbf{f}_Y } for MNIST classification

Layer	Description	Dimension	Filter dimension ($M \times N$)
\mathbf{x}	Input	784	—
\mathbf{f}	FC(sigmoid)	128	128×784
\mathbf{f}_Y	Output(softmax)	10	10×128

FC denotes the fully connected layer
 Dimension: the number of nodes (i.e., neurons) in every layer.
 Filter dimension ($M \times N$): a layer has M filters with $1 \times N$ dimension

A THE PARAMETERS OF DNNs CANNOT SATISFY THE I.I.D. PRIOR

A.1 A NEURAL NETWORK WITH A SINGLE HIDDEN LAYER ON THE MNIST DATASET

Based on the seminal work (Neal, 1994), we design a neural network with a single fully connected hidden layer (Figure 8) for classifying the benchmark MNIST dataset (LeCun et al., 1998). The input layer dimension is 1×784 , the hidden layer has 128 neurons with the sigmoid activation function, and the output layer is the softmax with 1×10 dimension. Table 2 summarizes the architecture of the network. In addition, $\{\alpha_{mn}\}_{m=1}^M$ denote the weights of the n th linear filter in the fully connected hidden layer and $\{\beta_{nl}\}_{n=1}^N$ are the weights of the l th linear filter in the output layer. Therefore, the hidden layer has 128 linear filters with 1×784 dimension, i.e., $\{\alpha_{mn}\}_{m=1}^{784}$. The output layer has totally 10 linear filters and their dimension is 1×128 , i.e., $\{\beta_{nl}\}_{n=1}^{128}$.




Figure 8: A neural network with a single hidden layer NN = { \mathbf{x} ; \mathbf{f} ; \mathbf{f}_Y }.

After the network is well trained (its training and testing accuracy is 97.4% and 96.9%, respectively), we calculate the absolute sample correlations, i.e., $|r(\alpha_n, \alpha_{n'})|$ and $|r(\beta_k, \beta_{k'})|$, and derive two absolute correlation matrices, i.e., $A'_{N \times N}$ and $B'_{L \times L}$, based on Equation 2. They are shown in Figure 9. We find that most $|r(\alpha_n, \alpha_{n'})|$ are close to zero, but $|r(\beta_k, \beta_{k'})|$ not. That indicates the parameters of the NN do not satisfy the necessary condition, thus they cannot satisfy the i.i.d. prior.




Figure 9: Two absolute correlation matrices for the weights in \mathbf{f} and \mathbf{f}_Y . In addition, 0.08 and 0.29 are the average of the absolute correlation coefficients for all the weights in \mathbf{f} and \mathbf{f}_Y , respectively.

Table 3: The architecture of the MLP for MNIST classification

Layer	Description	Dimension	Filter dimension ($M \times N$)
\mathbf{x}	Input	784	—
\mathbf{f}_1	FC(sigmoid)	128	128×784
\mathbf{f}_2	FC(sigmoid)	64	64×128
\mathbf{f}_Y	Output(softmax)	10	10×64

FC denotes the fully connected layer

Dimension: the number of nodes (i.e., neurons) in every layer.

Filter dimension ($M \times N$): a layer has M filters with $N \times 1$ dimension

A.2 A MLP ON THE MNIST DATASET

We design a MLP (Figure 1) to classify the benchmark MNIST dataset. The MLP has two hidden layers, each hidden layer is a fully connected layer with the sigmoid activation function, and the output layer is the softmax. Table 3 summarizes the architecture of the MLP.

After the MLP is well trained (its training and testing accuracy is 97.7% and 97.1%, respectively), we quantify the correlation of different weights in the same layer based on Equation 22 and derive the correlation matrix for all the weights in every hidden layer.

$$r(\boldsymbol{\alpha}_n, \boldsymbol{\alpha}_{n'}) = \frac{\sum_{m=1}^M (\alpha_{mn} - \bar{\alpha}_n)(\alpha_{mn'} - \bar{\alpha}_{n'})}{\sqrt{\sum_{m=1}^M (\alpha_{mn} - \bar{\alpha}_n)^2 \sum_{m=1}^M (\alpha_{mn'} - \bar{\alpha}_{n'})^2}} \approx \text{Corr}(\boldsymbol{\alpha}_n, \boldsymbol{\alpha}_{n'}) \quad (22)$$

where $\bar{\alpha}_n = \frac{1}{M} \sum_{m=1}^M \alpha_{mn}$ is the sample mean of $\{\alpha_{mn}\}_{m=1}^M$.

To examine if different weights are correlated, we only need to check whether their correlation coefficients are close to zero, so we take the absolute value of the correlation matrix element-wisely. Figure 10 visualizes the correlation matrices of each hidden layer, i.e., $A'_{N \times N}$, $B'_{K \times K}$, and $C'_{L \times L}$.

We find that most weights in \mathbf{f}_1 are uncorrelated because their correlation coefficients are very close to zero, i.e., $|r(\boldsymbol{\alpha}_n, \boldsymbol{\alpha}_{n'})| \approx 0$, but many weights in \mathbf{f}_2 and \mathbf{f}_Y are correlated to others, i.e., $|r(\boldsymbol{\beta}_k, \boldsymbol{\beta}_{k'})| \neq 0$ and $|r(\boldsymbol{\gamma}_l, \boldsymbol{\gamma}_{l'})| \neq 0$. In addition, we derive the average of the absolute correlation coefficients of all the weights in each layer and show them at the top of Figure 10, which further validates that the weights of some hidden layer are correlated.

Therefore, the parameters of the MLP cannot satisfy the i.i.d. prior.




Figure 10: Three absolute correlation matrices for the weights in \mathbf{f}_1 , \mathbf{f}_2 , and \mathbf{f}_Y . In addition, 0.08, 0.26, 0.31 are the average of the absolute correlation coefficients for all the weights in \mathbf{f}_1 , \mathbf{f}_2 , and \mathbf{f}_Y , respectively.

Table 4: The architectures of the CNN for CIFAR-10 classification

Layer	Description	Output dimension	Filters dimension	Correlation matrix
x	Input	$32 \times 32 \times 3$	—	—
f_1	Conv (3×3) + Maxpool	$16 \times 16 \times 64$	$3 \times 3 \times 3 \times 64$	$A_{192 \times 192}$
f_2	Conv (3×3) + Maxpool	$8 \times 8 \times 128$	$3 \times 3 \times 64 \times 128$	$B_{8192 \times 8192}$
f_3	FC + Sigmoid	1024	8192×1024	$C_{1024 \times 1024}$
f_4	FC + Sigmoid	256	1024×256	$D_{256 \times 256}$
f_Y	Output(softmax)	10	256×10	$E_{10 \times 10}$

A.3 A CNN ON THE CIFAR-10 DATASET

We design a CNN to classify the benchmark CIFAR-10 dataset (Krizhevsky, 2009). The CNN has two convolutional layers without activation function, two max-pooling layers for dimension reduction, and two fully connected layers with the sigmoid activation function. In addition, the output layer is the softmax. Table 4 summarizes the architecture of the CNN.

After we train the CNN well (its training and testing accuracy are 98.0% and 69.4%, respectively). We quantify the correlation of different weights in the same layer based on Equation 2, thereby deriving the absolute correlation matrix for all the weights in every hidden layer.

Since x has 3 channels and f_1 has 64 channels, f_1 has totally $3 \times 64 = 192$ filters. Hence, the dimension of the correlation matrix for f_1 is 192×192 . In particular, we flatten the output of the second convolutional layer, i.e., reshaping $8 \times 8 \times 128$ to 1×8192 , thus the filter dimension of f_3 is 8192×1024 . Table 4 shows the dimension of the correlation matrices for all the layers in the CNN.

Figure 11 shows the absolute correlation matrices for every layer. We can find that most weights are correlated in f_1 and f_2 , i.e., the weights in the convolutional layers are correlated. In addition, the weights in the output layer are also shown to be correlated. However, the weights in f_3 and f_4 are not correlated, because their correlation coefficients are close to zero. In other words, the weights in the fully connected layers are uncorrelated. Moreover, we derive the average of the absolute correlation coefficients of all the weights in each layer and show them at the top of Figure 11, which further validates the correlation of the weights in each layer. Overall, the parameters of the CNN cannot satisfy the necessary condition. Therefore, they cannot satisfy the i.i.d. prior.




Figure 11: Five absolute correlation matrices for the weights in the layer of CNN. The number at the top of the figure is the average of the absolute correlation coefficients for all the weights of the layer




Figure 12: Four synthetic images with different features. All images are sampled from $\mathcal{N}(0, 1)$. The only difference is that they are sorted in different diagonal directions by the ascending or descending orders. Image0 and Image1 are sorted in the primary diagonal direction by the ascending order and the descending order, respectively. Image2 and Image3 are sorted in the secondary diagonal direction by the ascending order and the descending order, respectively.

A.4 NEURAL NETWORKS WITH A SINGLE HIDDEN LAYER ON A SYNTHETIC DATASET

In order to further demonstrate that the parameters of DNNs cannot satisfy the i.i.d. assumption, we construct a neural network, i.e., $\text{NN} = \{\mathbf{x}; \mathbf{f}; \mathbf{f}_Y\}$ shown in Figure 8, to classify a synthetic dataset. The synthetic dataset consists of 4,800 32×32 grayscale images with 4 classes, and every synthetic image is sampled from the Gaussian distribution $\mathcal{N}(0, 1)$. To make the synthetic image have certain features that can be learned by neural networks, we sort the synthetic image in the primary or secondary diagonal directions by the ascending or descending orders. Thus, the synthetic dataset only has four features, i.e., four classes, and their corresponding images are drawn above.

There are two important reasons why we choose the synthetic dataset. Above all, compared to benchmark datasets containing too complex features, the synthetic dataset only has four simple features, thus we can easily show the correlation of different parameters by simply visualizing them. In addition, since a single synthetic image only has a single feature, we can easily examine the connection between the parameters of networks and the synthetic image.

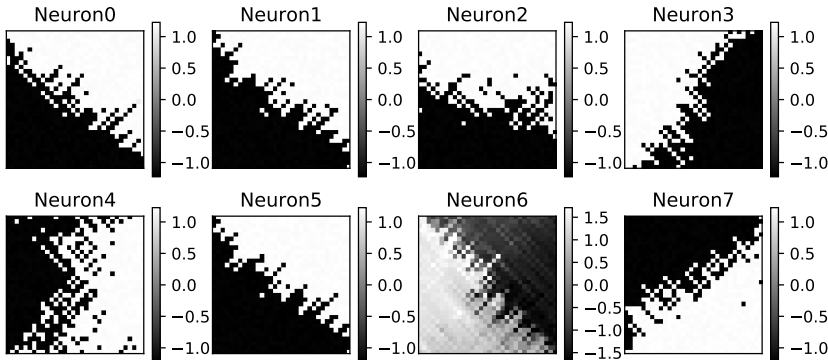


Figure 13: Eight neurons in the fully connected hidden layer. Their dimension is 32×32

Since the synthetic dataset is very simple, we set the hidden layer \mathbf{f} have only eight neurons and the output layer \mathbf{f}_Y have four output nodes. After training the network well, we show all the eight neurons in Figure 13. In contrast to the i.i.d. assumption, their weights show strong correlation in three aspects. First, weights show strong internal correlation within each neuron, i.e., if a weight has large magnitude, its neighbors have high probability to be large, and vice versa. Second, weights show strong external correlation between different neurons. For example, most weights of Neuron1 and Neuron5 at the same position show strong positive correlation, and most weights of Neuron3 and Neuron7 at the same position show strong negative correlation. Third, the weights show strong dependence on the training dataset. More specifically, since all the features of the synthetic dataset concentrate in the primary and secondary diagonal directions, most linear filters also demonstrate related features in the same directions.




Figure 14: The left figure visualizes the four linear filters in the output layer. The right figure visualizes the correlation matrix of the four filters in the output layer.

Figure 14 visualizes the four linear filters in the output layer. Though they not show too strong internal correlation because their dimension is too small, their weights at the same position still show strong external correlation, e.g., the weights of Filter0 and Filter1 at the same position show strong negative correlation. Moreover, we calculate the correlation coefficients of all the four filters based on Equation 2 and derive the correlation matrix show in Figure 14. It shows that the correlation coefficient of Filter0 and Filter1 is close to -1 , which further validates that they have strong negative correlation, hence the parameters of the output layer do not satisfy the i.i.d. prior.

Compared to the simulations using benchmark datasets, the synthetic dataset enable us to expose the correlation of different weights and their dependence on the training dataset in the more clearly way. Overall, we demonstrate that the parameters of neural networks are not i.i.d..

B PROOF: THE NECESSARY CONDITION FOR THE ACTIVATIONS OF HIDDEN LAYERS BEING INDEPENDENT




Figure 15: A MLP with two hidden layers, i.e., $\text{DNN} = \{\mathbf{x}; \mathbf{f}_1; \mathbf{f}_2; \mathbf{f}_Y\}$, in which \mathbf{f}_1 has N neurons, i.e., $\mathbf{f}_1 = \{f_{1n} = \sigma_1[g_{1n}(\mathbf{x})]\}_{n=1}^N$, where $g_{1n}(\mathbf{x}) = \sum_{m=1}^M \alpha_{mn} \cdot x_m + b_{1n}$ is the n th linear function with α_{mn} being the weight of the edge between x_m and f_{1n} and b_{1n} being the bias, and σ_1 is a non-linear activation function. In addition, $\mathbf{f}_2 = \{f_{2k} = \sigma_2[g_{2k}(\mathbf{f}_1)]\}_{k=1}^K$ where $g_{2k}(\mathbf{f}_1) = \sum_{n=1}^N \beta_{nk} \cdot f_{1n} + b_{2k}$ and $\mathbf{f}_Y = \{f_{Yl} = \sigma_3[g_{Yl}(\mathbf{f}_2)]\}_{l=1}^L$ where $g_{Yl}(\mathbf{f}_2) = \sum_{k=1}^K \gamma_{kl} \cdot f_{2k} + b_{Yl}$. They have the similar definitions as f_{1n} .

In the context of Bayesian probability, all the weights and the biases in the MLP (Figure 15) have prior distributions, i.e., $\alpha_{mn} \sim P(A_{mn})$, $\beta_{nk} \sim P(W_{nk})$, $b_{1n} \sim P(B_{1n})$, and $b_{2k} \sim P(B_{2k})$, thus we regard $G_{1n} = \sum_{m=1}^M A_{mn} X_m + B_{1n}$ as the random variable of $g_{1n}(\mathbf{x}) = \sum_{m=1}^M \alpha_{mn} \cdot x_m + b_{1n}$, and $G_{2k} = \sum_{n=1}^N W_{nk} F_{1n} + B_{2k}$ as the random variable of $g_{2k}(\mathbf{f}_1) = \sum_{n=1}^N \beta_{nk} \cdot f_{1n} + b_{2k}$. Therefore, the random variables of arbitrary activations in \mathbf{f}_1 and \mathbf{f}_2 , e.g., f_{1n} and f_{2k} , can be expressed as $F_{1n} = \sigma_1(G_{1n})$ and $F_{2k} = \sigma_2(G_{2k})$, respectively.

Based on the theorem (Appendix D) that functions of independent random variables are independent, we can derive that if $\{F_{2k}\}_{k=1}^K$ are independent and the activation function $\sigma_1(\cdot)$ is invertible, then $\{G_{2k}\}_{k=1}^K$ are independent. In other words, in order to prove $\{F_{2k}\}_{k=1}^K$ are independent, we only need to prove $\{G_{2k}\}_{k=1}^K$ being independent as long as $\sigma_2(\cdot)$ is invertible.

The necessary condition for $\{G_{2k}\}_{k=1}^K$ being independent is $\forall(k, k') \in S_1 = \{(k, k') \in \mathbb{Z}^2 | k \neq k', 1 \leq k \leq K, 1 \leq k' \leq K\}$, the covariance $Cov(G_{2k}, G_{2k'}) = 0$, which can be formulated as

$$Cov(G_{2k}, G_{2k'}) = Cov\left(\sum_{n=1}^N W_{nk}F_{1n} + B_{2k}, \sum_{n'=1}^N W_{n'k'}F_{1n'} + B_{2k'}\right) = 0 \quad (23)$$

Based on $Cov(X + Y, W + Z) = Cov(X + W) + Cov(X + Z) + Cov(Y + W) + Cov(Y + Z)$, we can derive that

$$\begin{aligned} Cov(G_{2k}, G_{2k'}) &= \sum_{n=1}^N \sum_{n'=1}^N Cov(W_{nk}F_{1n}, W_{n'k'}F_{1n'}) + \sum_{n=1}^N Cov(W_{nk}F_{1n}, B_{2k'}) \\ &\quad + \sum_{n'=1}^N Cov(W_{n'k'}F_{1n'}, B_{2k}) + Cov(B_{2k}, B_{2k'}) \end{aligned} \quad (24)$$

Based on the definition of covariance, i.e., $Cov(A, B) = E(AB) - E(A)E(B)$, we have

$$\begin{aligned} Cov(W_{nk}F_{1n}, W_{n'k'}F_{1n'}) &= E(W_{nk}W_{n'k'} \cdot F_{1n}F_{1n'}) - E(W_{nk}F_{1n})E(W_{n'k'}F_{1n'}) \\ Cov(W_{nk}F_{1n}, B_{2k'}) &= E(W_{nk}B_{2k'}F_{1n}) - E(W_{nk}F_{1n})E(B_{2k'}) \end{aligned} \quad (25)$$

Based on the law of total expectation, i.e, $E(AB) = E_A(AE_{B|A}(B))$, we have

$$\begin{aligned} E(W_{nk}W_{n'k'} \cdot F_{1n}F_{1n'}) &= E_{W_{nk}W_{n'k'}}(W_{nk}W_{n'k'} \cdot E_{F_{1n}F_{1n'}|W_{nk}W_{n'k'}}(F_{1n}F_{1n'})) \\ E(W_{nk}B_{2k'}F_{1n}) &= E_{W_{nk}B_{2k'}}(W_{nk}B_{2k'} \cdot E_{F_{1n}|W_{nk}B_{2k'}}(F_{1n})) \\ E(W_{nk}F_{1n}) &= E_{W_{nk}}(W_{nk} \cdot E_{F_{1n}|W_{nk}}(F_{1n})) \end{aligned} \quad (26)$$

Since the statistical properties of the input $\{F_{1n}\}_{n=1}^N$ are unrelated to the weights β_{nk} and biases b_{2k} , we have $E_{F_{1n}F_{1n'}|W_{nk}W_{n'k'}}(F_{1n}F_{1n'}) = E_{F_{1n}F_{1n'}}(F_{1n}F_{1n'})$, $E_{F_{1n}|W_{nk}B_{2k'}}(F_{1n}) = E_{F_{1n}}(F_{1n})$, and $E_{F_{1n}|W_{nk}}(F_{1n}) = E_{F_{1n}}(F_{1n})$. As a result, we can derive that

$$\begin{aligned} E(W_{nk}W_{n'k'} \cdot F_{1n}F_{1n'}) &= E_{W_{nk}W_{n'k'}}(W_{nk}W_{n'k'})E_{F_{1n}F_{1n'}}(F_{1n}F_{1n'}) \\ E(W_{nk}B_{2k'}F_{1n}) &= E_{W_{nk}B_{2k'}}(W_{nk}B_{2k'})E_{F_{1n}}(F_{1n}) \\ E(W_{nk}F_{1n}) &= E_{W_{nk}}(W_{nk})E_{F_{1n}}(F_{1n}) \end{aligned} \quad (27)$$

Substituting Equation (27) for the corresponding expectations in Equation (25), we have

$$\begin{aligned} Cov(W_{nk}F_{1n}, W_{n'k'}F_{1n'}) &= E(W_{nk}W_{n'k'})E(F_{1n}F_{1n'}) - E(W_{nk})E(W_{n'k'})E(F_{1n})E(F_{1n'}) \\ Cov(W_{nk}F_{1n}, B_{2k'}) &= E(W_{nk}B_{2k'})E(F_{1n}) - E(W_{nk})E(B_{2k'})E(F_{1n}) \end{aligned} \quad (28)$$

Assuming $\{F_{1n}\}_{n=1}^N$ are i.i.d. and $F_{1n} \sim P(F_1)$, we have $E(F_{1n}F_{1n'}) = E(F_{1n})E(F_{1n'}) = E^2(F_1)$. Hence, we can derive that

$$\begin{aligned} Cov(W_{nk}F_{1n}, W_{n'k'}F_{1n'}) &= Cov(W_{nk}, W_{n'k'})E^2(F_1) \\ Cov(W_{nk}F_{1n}, B_{2k'}) &= Cov(W_{nk}, B_{2k'})E(F_1) \end{aligned} \quad (29)$$

Substituting Equation (29) for the corresponding covariances in Equation (24), we have

$$\begin{aligned} Cov(G_{2k}, G_{2k'}) &= \sum_{n=1}^N \sum_{n'=1}^N Cov(W_{nk}, W_{n'k'})E^2(F_1) + \sum_{n=1}^N Cov(W_{nk}, B_{2k'})E(F_1) \\ &\quad + \sum_{n'=1}^N Cov(W_{n'k'}, B_{2k})E(F_1) + Cov(B_{2k}, B_{2k'}) \end{aligned} \quad (30)$$

Based on $Cov(X + Y, W + Z) = Cov(X + W) + Cov(X + Z) + Cov(Y + W) + Cov(Y + Z)$, we can derive that

$$\begin{aligned} Cov(G_{2k}, G_{2k'}) &= Cov\left(\sum_{n=1}^N W_{nk}, \sum_{n'=1}^N W_{n'k'}\right)E^2(F_1) + Cov\left(\sum_{n=1}^N W_{nk}, B_{2k'}\right)E(F_1) \\ &\quad + Cov\left(\sum_{n'=1}^N W_{n'k'}, B_{2k}\right)E(F_1) + Cov(B_{2k}, B_{2k'}) \end{aligned} \quad (31)$$

Therefore, given $\sigma_2(\cdot)$ is invertible and $\{F_{1n}\}_{n=1}^N$ are i.i.d., the necessary condition for $\{F_{2k}\}_{k=1}^K$ being independent can be expressed as

$$\begin{aligned} Cov(G_{2k}, G_{2k'}) &= Cov(W_k, W_{k'})E^2(F_1) + Cov(W_k, B_{2k'})E(F_1) \\ &\quad + Cov(W_{k'}, B_{2k})E(F_1) + Cov(B_{2k}, B_{2k'}) = 0 \end{aligned} \quad (32)$$

where $W_k = \sum_{n=1}^N W_{nk}$ and $W_{k'} = \sum_{n=1}^N W_{n'k'}$.

Overall, in order to guarantee the activations of a hidden layer are independent with each other in the context of Bayesian probability, the prior distributions of the weights and biases of the hidden layer should satisfy the above necessary condition given the inputs of the hidden layer are i.i.d. and the activation functions are invertible.

C PROOF: THE NECESSARY CONDITION FOR THE ACTIVATIONS OF HIDDEN LAYERS BEING IDENTICALLY DISTRIBUTED

Given the same DNN drawn in Figure 15, we assume $\sigma_2(\cdot)$ being strictly increasing and differentiable, so $\sigma_2(\cdot)$ is invertible and its inverse $\sigma_2^{-1}(\cdot)$ is also strictly increasing. The cumulative distribution function of F_{2k} can be expressed as

$$\begin{aligned} \Phi_{F_{2k}}(f) &= \phi(F_{2k} \leq f) \\ &= \phi(\sigma_1(G_{2k}) \leq f) \\ &= \phi(G_{2k} \leq \sigma_1^{-1}(f)) \\ &= \Phi_{G_{2k}}(\sigma_1^{-1}(f)) \end{aligned} \quad (33)$$

where $\phi(F_{2k} \leq f)$ is the probability of F_{2k} takes on a value less than or equal to f . Subsequently, we can obtain that

$$\begin{aligned} P_{F_{2k}}(f) &= \frac{\partial \Phi_{F_{2k}}(f)}{\partial f} = \frac{\partial \Phi_{G_{2k}}(\sigma_1^{-1}(f))}{\partial f} \\ &= P_{G_{2k}}(\sigma_1^{-1}(f)) \frac{\partial \sigma_1^{-1}(f)}{\partial f} \end{aligned} \quad (34)$$

Equation (34) means that if the activation function $\sigma_2(\cdot)$ is strictly increasing and differentiable, $\{F_{2k}\}_{k=1}^K$ being identically distributed implies $\{G_{2k}\}_{k=1}^K$ being identically distributed.

In the context of Bayesian probability, $G_{2k} = \sum_{n=1}^N W_{nk}F_{1n} + B_{2k}$, we can derive that

$$E(G_{2k}) = \sum_{n=1}^N E(W_{nk}F_{1n}) + E(B_{2k}) \quad (35)$$

Based on the law of total expectation and Equation 27, we have

$$E(W_{nk}F_{1n}) = E_{W_{nk}}(W_{nk})E_{F_{1n}}(F_{1n}) \quad (36)$$

Assuming $\{F_{1n}\}_{n=1}^N$ are i.i.d. and $F_{1n} \sim F_1$, we have $E(F_{1n}) = E(F_1)$. Hence, we can derive that

$$E(G_{2k}) = E(F_1) \sum_{n=1}^N E(W_{nk}) + E(B_{2k}) \quad (37)$$

Based on the property of expectation, i.e., $E(A) + E(B) = E(A + B)$, we have

$$E(G_{2k}) = E(F_1)E\left(\sum_{n=1}^N W_{nk}\right) + E(B_{2k}) \quad (38)$$

Therefore, given $\sigma_2(\cdot)$ is strictly increasing and differentiable and $\{F_{1n}\}_{n=1}^N$ are i.i.d., the necessary condition for $\{F_{2k}\}_{k=1}^K$ being identically distributed is that $\forall(n, n') \in S_1$, we have

$$E(F_1)E(W_k) + E(B_{2k}) = E(F_1)E(W_{k'}) + E(B_{2k'}) \quad (39)$$

where $W_k = \sum_{n=1}^N W_{nk}$ and $W_{k'} = \sum_{n=1}^N W_{nk'}$.

Overall, in order to guarantee the activations of a hidden layer are identically distributed, the prior distributions of the weights and the biases of the hidden layer should satisfy the above necessary condition in the context of Bayesian probability given the inputs of the hidden layer are i.i.d. and the activation functions are strictly increasing and differentiable.

D THEOREM: FUNCTIONS OF INDEPENDENT RANDOM VARIABLES ARE INDEPENDENT

Let \mathbf{X} and \mathbf{Y} be independent random variables on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let g and h be real-valued functions defined on the codomains of \mathbf{X} and \mathbf{Y} , respectively. Then $g(\mathbf{X})$ and $h(\mathbf{Y})$ are independent random variables.

Proof: Let $\mathbf{A} \subseteq \mathbb{R}$ and $\mathbf{B} \subseteq \mathbb{R}$ be the range of g and h , the joint distribution between $g(\mathbf{X})$ and $h(\mathbf{Y})$ can be formulate as $\mathbb{P}(g(\mathbf{X}) \in \mathbf{A}, h(\mathbf{Y}) \in \mathbf{B})$. Let $g^{-1}(\mathbf{A})$ and $h^{-1}(\mathbf{B})$ denote the preimages of \mathbf{A} and \mathbf{B} , respectively, we have

$$\mathbb{P}(g(\mathbf{X}) \in \mathbf{A}, h(\mathbf{Y}) \in \mathbf{B}) = \mathbb{P}(\mathbf{X} \in g^{-1}(\mathbf{A}), \mathbf{Y} \in h^{-1}(\mathbf{B})) \quad (40)$$

Based on the definition of independence, we can derive that

$$\begin{aligned} \mathbb{P}(g(\mathbf{X}) \in \mathbf{A}, h(\mathbf{Y}) \in \mathbf{B}) &= \mathbb{P}(\mathbf{X} \in g^{-1}(\mathbf{A}))\mathbb{P}(\mathbf{Y} \in h^{-1}(\mathbf{B})) \\ &= \mathbb{P}(g(\mathbf{X}) \in \mathbf{A})\mathbb{P}(h(\mathbf{Y}) \in \mathbf{B}) \end{aligned} \quad (41)$$

Based on the definition of preimage, we can derive that

$$\mathbb{P}(g(\mathbf{X}) \in \mathbf{A}, h(\mathbf{Y}) \in \mathbf{B}) = \mathbb{P}(g(\mathbf{X}) \in \mathbf{A})\mathbb{P}(h(\mathbf{Y}) \in \mathbf{B}) \quad (42)$$

Therefore, $g(\mathbf{X})$ and $h(\mathbf{Y})$ are independent random variables.

E THE NECESSARY CONDITIONS FOR ACTIVATIONS BEING I.I.D. ARE VALID FOR CONVOLUTIONAL LAYERS

Assuming a DNN = $\{\mathbf{x}; \mathbf{f}_1; \dots; \mathbf{f}_I; \mathbf{f}_Y\}$ has two adjacent convolutional layers \mathbf{f}_i and \mathbf{f}_{i+1} . The first convolutional layer has N channels, i.e. $\mathbf{f}_i = \{\mathbf{f}_i^n\}_{n=1}^N$, and the second convolutional layer has K channels, i.e. $\mathbf{f}_{i+1} = \{\mathbf{f}_{i+1}^k\}_{k=1}^K$. Therefore, the dimension of the convolutional filters \mathbf{S}_{i+1} in \mathbf{f}_{i+1} are $W \times H \times N \times K$, i.e., the dimension of a single filter connecting the \mathbf{f}_i^n channel and the \mathbf{f}_{i+1}^k channel are $D(\mathbf{S}_{i+1}^{(k,n)}) = W \times H$. As a result, the \mathbf{f}_{i+1}^k channel can be formulated as

$$\mathbf{f}_{i+1}^k = \sigma_{i+1}\left(\sum_{n=1}^N \mathbf{S}_{i+1}^{(k,n)} \circ \mathbf{f}_i^n\right) + b_{i+1}^k \cdot \mathbf{1} \quad (43)$$

where σ_{i+1} is the activation function and b_{i+1} is the bias for the $(i + 1)$ th channel.

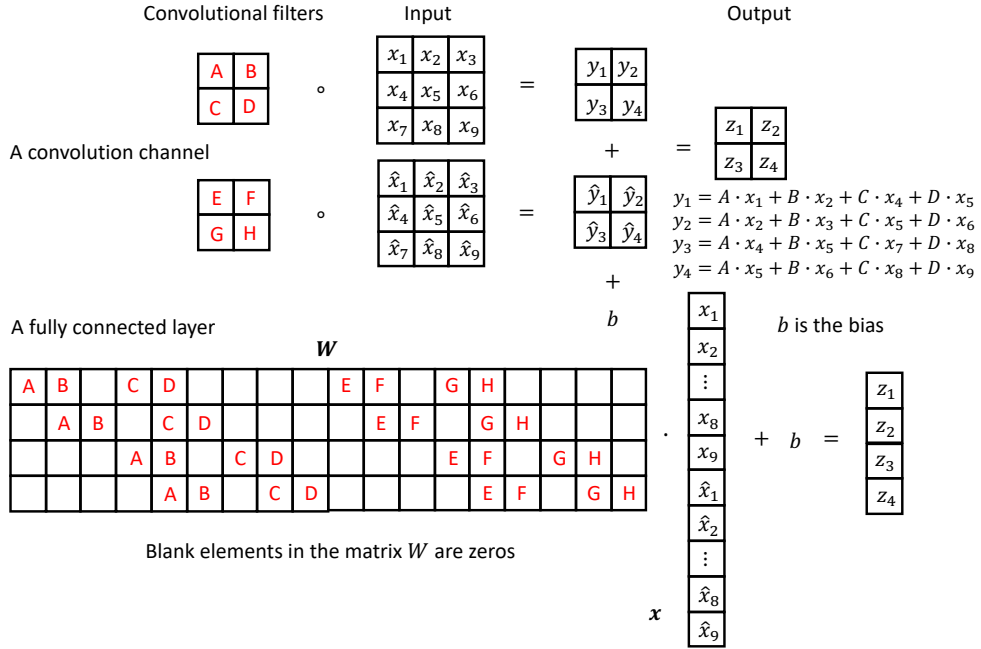


Figure 16: The equivalence between a convolution channel and a fully connected layer.

Since Garriga-Alonso et al. (2018) demonstrate the equivalence between a convolution operation and the matrix multiplication, we can regard the convolutional channel f_{i+1}^k as a fully connected layer. Specifically, the input of the virtual fully connected layer is the same as the input of the convolutional channel, the output of the virtual fully connected layer is the vectorized output of the convolutional channel, and the weights of the virtual fully connected layer depend on the convolutional filter and the spatial location of the output of the convolutional channel.

Without considering the activation function, the equivalence between a convolutional channel and a fully connected layer can be explained in Figure 16, in which the first convolutional layer has two channels, i.e., $f_i = \{x; \hat{x}\}$, where $x = \{x_1, \dots, x_9\}$ and $\hat{x} = \{\hat{x}_1, \dots, \hat{x}_9\}$, and the k th channel of the second convolutional layer is $f_{i+1}^k = \{z_1, z_2, z_3, z_4\}$. In addition, the dimension of the convolutional filter S_{i+1} is $2 \times 2 \times 2 \times 1$, where $S_{i+1}^{(k,1)} = [A, B; C, D]$ and $S_{i+1}^{(k,2)} = [E, F; G, H]$. Figure 16 shows that the output of f_{i+1}^k can be expressed as a fully connected layer. The input of the fully connected layer is the vector x , the weights of each neuron is represented by each row of the matrix W , and the bias is b , thus $f_{i+1}^k = W \cdot x + b \cdot \mathbf{1}$.

If we consider the activation function and make a precise statement, f_{i+1}^k can be formulated as

$$f_{i+1}^k = \sigma_{i+1}(W_{i+1}^k \cdot f_i + b_{i+1}^k \cdot \mathbf{1}) \quad (44)$$

where W_{i+1}^k is the matrix corresponding to all the convolutional filters in the k th channel.

To derive the necessary conditions for two arbitrary different channels are i.i.d., we need to derive the necessary conditions for two arbitrary activations in two arbitrary different channels are i.i.d. For example, if the m th activation of the k channel and the m' th activation of the k' channel in f_{i+1} are i.i.d., the necessary conditions are $Cov(f_{i+1}^k(m), f_{i+1}^{k'}(m')) = 0$, $E(f_{i+1}^k(m)) = E(f_{i+1}^{k'}(m'))$.

Since we demonstrate the equivalence between a convolutional channel and a fully connected layer, the necessary conditions for activations being i.i.d. in a fully connected layer are also valid for a convolutional layer. In other words, we can also use the Equation 3, Equation 4, and Equation 5 to examine whether activations being i.i.d. is valid in a convolutional layer under the assumption that all activations being i.i.d. in the previous layer.

Table 5: The statistical measures for examining the necessary conditions

Layer	Activation expectation	Covariance	Weight expectation
\mathbf{f}_1	$E(X)$	$Cov(A_n, A_{n'})$	$E(A_n)$
\mathbf{f}_2	$E(F_1)$	$Cov(W_k, W_{k'})$	$E(W_k)$
\mathbf{f}_Y	$E(F_2)$	$Cov(C_l, C_{l'})$	$E(C_l)$

$$\{x_m\}_{m=1}^M \sim P(X), \{f_{1n}\}_{n=1}^N \sim P(F_1), \{f_{2k}\}_{k=1}^K \sim P(F_2)$$

$$A_n = \sum_{m=1}^M A_{mn}, \text{ where } \alpha_{mn} \sim P(A_{mn})$$

$$W_k = \sum_{n=1}^N W_{nk}, \text{ where } \beta_{nk} \sim P(W_{nk})$$

$$C_l = \sum_{k=1}^K C_{kl}, \text{ where } \gamma_{kl} \sim P(C_{kl})$$

F ACTIVATIONS BEING I.I.D. CANNOT BE VALID FOR ALL THE HIDDEN LAYERS OF DNNs

In this section, we demonstrate that activations being i.i.d. cannot be valid for all the hidden layers of DNNs through showing that the hidden layers of two typical DNNs, i.e, the MLP and the CNN, cannot satisfy the necessary conditions for activations being i.i.d..

F.1 THE MLP ON THE MNIST DATASET

The architecture of the MLP for classifying the MNIST dataset is visualized in Figure 15, and its specific parameters are summarized in Table 3. Since the activation functions of all the hidden layers in the MLP are the sigmoid function, which is strictly increasing, differentiable, and invertible, we only need to examine the first two necessary conditions in the MLP.

Since we restrict all the hidden layers from using the bias, the first two necessary conditions for activations being i.i.d. can be simplified as

$$\forall (k, k') \in S_1 = \{(k, k') \in \mathbb{Z}^2 | k \neq k', 1 \leq k \leq K, 1 \leq k' \leq K\}, \text{ we must have} \quad (45)$$

$$Cov(G_{2k}, G_{2k'}) = Cov(W_k, W_{k'})E^2(F_1) = 0$$

$$E(F_1)E(W_k) = E(F_1)E(W_{k'}) \quad (46)$$

Based on the architecture of the MLP, Table 5 summarizes all the statistical measures for each layer of the MLP. Therefore, the key to examine if the layers of the MLP satisfy the necessary conditions is to estimate the statistical measures.

To estimate the expectations $E(X)$, $E(F_1)$, and $E(F_2)$, we use the sample means \bar{x} , $\bar{\mathbf{f}}_1$, and $\bar{\mathbf{f}}_2$, respectively. The definition of the sample mean is

$$\bar{x} = \frac{1}{M} \sum_{m=1}^M x_m, \bar{\mathbf{f}}_1 = \frac{1}{N} \sum_{n=1}^N f_2(n), \bar{\mathbf{f}}_2 = \frac{1}{K} \sum_{k=1}^K f_2(k) \quad (47)$$

After the MLP is well-trained, we derive the sample means as $\bar{x} = 0.131$, $\bar{\mathbf{f}}_1 = 0.436$, and $\bar{\mathbf{f}}_2 = 0.498$ given the benchmark MNIST testing dataset.

Taking into account the estimation error, we can regard $E(X) = 0$. As a result, we can derive that $\forall (n, n'), Cov(G_{1n}, G_{1n'}) = 0$ and $E(G_{1n}) = E(G_{1n'}) = 0$, thus \mathbf{f}_1 satisfy the necessary conditions for activations being i.i.d..

However, $\bar{\mathbf{f}}_1 = 0.436$ and $\bar{\mathbf{f}}_2 = 0.498$ imply that $E(F_1) \neq 0$ and $E(F_2) \neq 0$ even considering the estimation error. Therefore, we need to estimate $Cov(W_k, W_{k'})$ and $E(W_k)$ to examine whether the activations $\{f_{2k}\}_{k=1}^K$ being i.i.d..

Since the training dataset is commonly viewed as i.i.d., we can use the same training dataset to train the MLP several times to derive multiple independent observation samples of $\{\beta_{nk}(t)\}_{t=1}^T$ to estimate the random variable W_{nk} , thereby estimating $Cov(W_k, W_{k'})$ and $E(W_k)$.




Figure 17: The expectation of the weights summation, i.e., $E(A_n)$, $E(W_k)$, and $E(C_l)$, in each layer. The number at the top of the figure indicates the variance of the points drawn in the figure.

More specifically, we use the sample correlation $r(\beta_k, \beta_{k'})$ to check whether $Cov(W_k, W_{k'})$ are close to zero and use the sample mean $\bar{\beta}_k$ to estimate $E(W_k)$.

$$r(\beta_k, \beta_{k'}) = \frac{\sum_{t=1}^T (\beta_k(t) - \bar{\beta}_k)(w_{k'}(t) - \bar{\beta}_{k'})}{\sqrt{\sum_{t=1}^T (\beta_k(t) - \bar{\beta}_k)^2 \sum_{t=1}^T (\beta_{k'}(t) - \bar{\beta}_{k'})^2}} \approx Corr(W_k, W_{k'}) \quad (48)$$

where $\bar{\beta}_k = \frac{1}{T} \sum_{t=1}^T \beta_k(t)$ is the sample mean of $\{\beta_k(t)\}_{t=1}^T$ and $\beta_k(t) = \sum_{n=1}^N \beta_{nk}(t)$.

In order to make a tradeoff between the estimation precision and computation complexity, we use the same training dataset to train the MLP 20 times to derive the samples, i.e., $\{\beta_{nk}(t)\}_{t=1}^{20}$. For an activation f_{2k} , it has 128 weights because its previous layer has 128 neurons. Hence, a single sample $\beta_k(t) = \sum_{n=1}^{128} \beta_{nk}(t)$. Based on the 20 samples, i.e., $\{\beta_k(t)\}_{t=1}^{20}$, we can calculate $r(\beta_k, \beta_{k'})$ to examine whether $Cov(W_k, W_{k'})$ are close to zero and derive the sample mean to estimate $E(W_k)$. We use the same method to estimate $Cov(A_n, A_{n'})$, $E(A_n)$, $Cov(C_l, C_{l'})$, and $E(C_l)$.

Finally, we shows the absolute correlation matrices i.e., $A_{N \times N}$, $W_{K \times K}$, and $C_{L \times L}$, in Figure 3. Since many elements in the three matrices are far from zero, we conclude that their corresponding covariances, i.e., $Cov(A_n, A_{n'})$, $Cov(W_k, W_{k'})$, and $Cov(C_k, C_{k'})$, are not close to zero based on connection between correlation and covariance

$$Corr(W_k, W_{k'}) = \frac{Cov(W_k, W_{k'})}{\sigma_{W_k} \sigma_{W_{k'}}} \quad (49)$$

where $\sigma_{W_k} > 0$ is the standard derivation of W_k .

Figure 17 visualizes the expectations, i.e., $E(A_n)$, $E(W_k)$, and $E(C_l)$, for each layer in the MLP. We find that some expectations show great variations, especially $\{E(A_n)\}_{n=1}^N$ and $\{E(C_l)\}_{l=1}^L$. In other words, the MLP cannot guarantee that all the expectations of the weights summation are equivalent even though we take into account the estimation error. Therefore, the MLP cannot satisfy the necessary condition for activations being identically distributed.

Overall, we conclude that not all the hidden layers of the MLP can satisfy the necessary conditions for activations being i.i.d. under the assumption that the activations of the previous layer are i.i.d.

F.2 THE CNN ON THE CIFAR-10 DATASET

The architecture of the CNN for classifying the CIFAR-10 dataset is in Table 4. In particular, the two convolutional layers do not have activations function, i.e., their activation functions can be viewed as $y = x$, which is strictly increasing, differentiable, and invertible. The sigmoid activation function is also strictly increasing, differentiable, and invertible, thus we only need to examine the first two necessary conditions in the CNN. To simplify derivation and decrease computation complexity, we assume the max-pooling layers are special activation functions and satisfy the third necessary condition for activations being i.i.d.

Table 6: The statistical measures for examining the necessary conditions of each layer in the CNN

	f_1	f_2	f_3	f_4	f_Y
Description	Conv Maxpool	Conv Maxpool	FC Sigmoid	FC Sigmoid	Output Softmax
Activation Dim.	$16 \times 16 \times 64$	$8 \times 8 \times 128$	1024×1	256×1	10×1
Filter Dim.	$3 \times 3 \times 3 \times 64$	$3 \times 3 \times 64 \times 128$	8192×1024	1024×256	256×10
Weight Dim.	$M \times N$ (27×64)	$N \times K$ (576×128)	$K \times L$ (8192×1024)	$L \times Q$ (1024×256)	$Q \times R$ (256×10)
Weight PDF	$a_{mn} \sim P(A_{mn})$	$b_{nk} \sim P(B_{nk})$	$c_{kl} \sim P(C_{kl})$	$d_{lq} \sim P(D_{lq})$	$f_{qr} \sim P(F_{qr})$
Sum. PDF	$A_n = \sum_{m=1}^M A_{mn}$	$B_k = \sum_{n=1}^N B_{nk}$	$C_l = \sum_{k=1}^K C_{kl}$	$D_q = \sum_{l=1}^L D_{lq}$	$F_r = \sum_{q=1}^Q F_{qr}$
Covariance	$Cov(A_n, A_{n'})$	$Cov(B_k, B_{k'})$	$Cov(C_l, C_{l'})$	$Cov(D_q, D_{q'})$	$Cov(F_r, F_{r'})$
Sum. Exp.	$E(A_n)$	$E(B_k)$	$E(C_l)$	$E(D_q)$	$E(F_r)$
Activation Exp.	$E(X)$	$E(F_1)$	$E(F_2)$	$E(F_3)$	$E(F_4)$

Dim. is short for dimension. Sum. is short for Weights Summation. Exp. is short for Expectation.

Similar as the MLP, we also restrict all the layers of the CNN from using the bias, thus the first two necessary conditions for activations being i.i.d. can be simplified as

$$\forall (k, k') \in S_1 = \{(k, k') \in \mathbb{Z}^2 | k \neq k', 1 \leq k \leq K, 1 \leq k' \leq K\}, \text{ we must have} \quad (50)$$

$$Cov(G_{2k}, G_{2k'}) = Cov(W_k, W_{k'})E^2(F_1) = 0$$

$$E(F_1)E(W_k) = E(F_1)E(W_{k'}) \quad (51)$$

Table 6 summarizes all the statistical measures for every layer of the CNN based on the architecture of the CNN in Table 4. The definitions of most random variables are similar as the MLP except the dimension of the weights. To check if two arbitrary convolutional channels are i.i.d., we need to know all the weights of two convolutional channels. Based on the equivalence between a convolutional channel and a fully connected layer (Appendix E), all the weights of a single convolutional channel consist of all the convolutional filters related to the convolutional channel. For example, the dimension of all the convolutional filters related to a convolutional channel in f_1 are $3 \times 3 \times 3$, thus a single convolutional channel in f_1 has 27 weights, so the weights dimension of f_1 is 27×64 .

After training the CNN well (its training accuracy is 96.2% and testing accuracy is 67.9%), we use the same methods as the MLP to estimate the above statistical measures in Table 6.




Figure 18: Five absolute correlation matrices for the weight summations in each layer. The number at the top is the average of the absolute correlation coefficients for all the weights summations of the layer




Figure 19: The expectation of the weights summation of each layer in the CNN. The number at the top of the figure indicates the variance of the points drawn in the figure..

For the activation expectations of each layer in the CNN, we have $\bar{x} = -0.023$, $\bar{f_1} = -707.074$, $\bar{f_2} = -6347.836$, $\bar{f_3} = 36.870$, $\bar{f_4} = 15.131$, thus all the expectations cannot be zero except the input layer, i.e., $E(X) = 0$, $E(F_1) \neq 0$, $E(F_2) \neq 0$, $E(F_3) \neq 0$, $E(F_4) \neq 0$.

For the covariances, we show all the correlation matrices of each layer of the CNN in Figure 18. We find that many correlation coefficients are far from zero in all the correlation matrices, thus their corresponding covariances are not close zero either. Based on the estimations for the activation expectation and the covariance between different weights summations in each layer, we can derive that activations being independent cannot be valid for all the layers of the CNN.

For the weights summation expectation, we visualize all the expectations for each layer in the CNN in Figure 19. We find that some expectations have great variations, especially $\{E(C_l)\}_{l=1}^L$ and $\{E(D_q)\}_{q=1}^Q$, which means that most expectations of weights summations of different neurons in f_3 and f_4 cannot be equivalent. As a result, the activations of the two fully connected layers in the CNN cannot be identically distributed.

Overall, we conclude that not all the hidden layers of the CNN can satisfy the necessary conditions for activations being i.i.d. under the assumption that the activations of the previous layer are i.i.d.

G THE EQUIVALENCE BETWEEN THE GRADIENT DESCENT ALGORITHM AND THE FIRST ORDER APPROXIMATION

In the context of deep learning, most learning algorithms belong to the gradient descent algorithm. Given a DNN = $\{\mathbf{x}; \mathbf{f}_1; \dots; \mathbf{f}_I; \mathbf{f}_Y\}$ and a cost function $H[\mathbf{f}_Y, P(\mathbf{Y}|\mathbf{X})]$, where $P(\mathbf{Y}|\mathbf{X})$ is the true distribution of the training label given the corresponding training dataset. Let θ be the parameters of the DNN, the gradient descent aims to optimize θ by minimizing $H[\mathbf{f}_Y, P(\mathbf{Y}|\mathbf{X})]$ (Rumelhart et al., 1986). We typically update θ iteratively to derive θ^* , which can be expressed as

$$\theta_{t+1} = \theta_t - \alpha \nabla_{\theta_t} H[\mathbf{f}_Y, P(\mathbf{Y}|\mathbf{X})] \quad (52)$$

where $\nabla_{\theta_t} H[\mathbf{f}_Y, P(\mathbf{Y}|\mathbf{X})]$ is the Jacobian matrix of $H[\mathbf{f}_Y, P(\mathbf{Y}|\mathbf{X})]$ with respect to θ_t at the t th iteration, and $\alpha > 0$ denotes a constant indicating the learning rate. Since $P(\mathbf{Y}|\mathbf{X})$ is constant, we denote $H(\mathbf{f}_Y)$ as $H[\mathbf{f}_Y, P(\mathbf{Y}|\mathbf{X})]$ for simplifying the following derivation.

Since the functions of all hidden layers are differentiable and the output of a hidden layer is the input of its next layer, the Jacobian matrix of $H(\mathbf{f}_Y)$ with respect to the parameters of the i th hidden layer, i.e., $\nabla_{\theta^{(i)}} H(\mathbf{f}_Y)$, can be expressed as follows based on the chain rule.

$$\nabla_{\theta^{(i)}} H(\mathbf{f}_Y) = \nabla_{\mathbf{f}_Y} H(\mathbf{f}_Y) \cdot \nabla_{\mathbf{f}_I} \mathbf{f}_Y \cdot \prod_{j=i+1}^I \nabla_{\mathbf{f}_{j-1}} \mathbf{f}_j \cdot \nabla_{\theta^{(i)}} \mathbf{f}_i \quad (53)$$

where $\theta^{(i)}$ denote the parameters of the i th hidden layer.

Based on Equation 52 and Equation 53, $\theta^{(i)}$ can be learned by the gradient descent method as

$$\theta_{t+1}^{(i)} = \theta_t^{(i)} - \alpha [\nabla_{\theta_t^{(i)}} H(\mathbf{f}_Y)] \quad (54)$$

Table 7 summarizes the backpropagation training procedure for the MLP in Figure 15.

Table 7: One iteration of the backpropagation training procedure for the MLP

Layer	Gradients update $\nabla_{\theta_{t(i)}} H(\mathbf{f}_Y)$	Parameters and activations update
\mathbf{f}_Y	$\nabla_{\mathbf{f}_Y} H(\mathbf{f}_Y) \nabla_{\theta_t(Y)} \mathbf{f}_Y$	\downarrow $\theta_{t+1}(Y) = \theta_t(Y) - \alpha [\nabla_{\theta_t(Y)} H(\mathbf{f}_Y)]$, $\mathbf{f}_Y(\mathbf{f}_2, \theta_{t+1}(Y))$
\mathbf{f}_2	$\nabla_{\mathbf{f}_Y} H(\mathbf{f}_Y) \nabla_{\mathbf{f}_2} \mathbf{f}_Y \nabla_{\theta_t(2)} \mathbf{f}_2$	\downarrow $\theta_{t+1}(2) = \theta_t(2) - \alpha [\nabla_{\theta_t(2)} H(\mathbf{f}_Y)]$, $\mathbf{f}_2(\mathbf{f}_1, \theta_{t+1}(2))$
\mathbf{f}_1	$\nabla_{\mathbf{f}_Y} H(\mathbf{f}_Y) \nabla_{\mathbf{f}_2} \mathbf{f}_Y \nabla_{\mathbf{f}_1} \mathbf{f}_2 \nabla_{\theta_t(1)} \mathbf{f}_1$	\downarrow $\theta_{t+1}(1) = \theta_t(1) - \alpha [\nabla_{\theta_t(1)} H(\mathbf{f}_Y)]$, $\mathbf{f}_1(\mathbf{x}, \theta_{t+1}(1))$
\mathbf{x}	—	—

The uparrow and downarrow indicate the order of gradients and parameters(activations) update, respectively.

If an arbitrary function \mathbf{f} is differentiable at point \mathbf{p}^* in \mathbb{R}^N and its differential is represented by the Jacobian matrix $\nabla_{\mathbf{p}} \mathbf{f}$, the first order approximation of \mathbf{f} near the point \mathbf{p} can be formulated as

$$\mathbf{f}(\mathbf{p}) - \mathbf{f}(\mathbf{p}^*) = (\nabla_{\mathbf{p}^*} \mathbf{f}) \cdot (\mathbf{p} - \mathbf{p}^*) + o(\|\mathbf{p} - \mathbf{p}^*\|) \quad (55)$$

where $o(\|\mathbf{p} - \mathbf{p}^*\|)$ is a quantity that approaches zero much faster than $\|\mathbf{p} - \mathbf{p}^*\|$ approaches zero.

Based on the equivalence between the gradient descent method and the first order approximation (Battiti, 1992), updating the activations of the hidden layers of the MLP in Figure 15 during the backpropagation training procedure can be approximated as

$$\begin{aligned} \mathbf{f}_2[\mathbf{f}_1, \theta_{t+1}(2)] &\approx \mathbf{f}_2[\mathbf{f}_1, \theta_t(2)] + (\nabla_{\theta_t(2)} \mathbf{f}_2) \cdot [\theta_{t+1}(2) - \theta_t(2)] \\ \mathbf{f}_1[\mathbf{x}, \theta_{t+1}(1)] &\approx \mathbf{f}_1[\mathbf{x}, \theta_t(1)] + (\nabla_{\theta_t(1)} \mathbf{f}_1) \cdot [\theta_{t+1}(1) - \theta_t(1)] \end{aligned} \quad (56)$$

where $\mathbf{f}_2[\mathbf{f}_1, \theta_t(2)]$ denote the activations of the second hidden layer based on the parameters learned in the t th iteration, i.e., $\theta_t(2)$, given the activations of the first hidden layer, i.e., \mathbf{f}_1 . The definitions of $\mathbf{f}_2[\mathbf{f}_1, \theta_{t+1}(2)]$ and $\mathbf{f}_1[\mathbf{x}, \theta_{t+1}(1)]$ are the same as $\mathbf{f}_2[\mathbf{f}_1, \theta_t(2)]$.

Because \mathbf{f}_2 has K neurons, i.e., $\mathbf{f}_2 = \{f_{2k} = \sigma_2(\sum_{n=1}^N \beta_{nk} \cdot f_{1n} + b_{2k})\}_{k=1}^K$, and each neuron has $N + 1$ parameters, i.e., $\theta(2) = \{[\beta_{1k}; \dots; \beta_{Nk}; b_{2k}]\}_{k=1}^K$, the dimension of $\nabla_{\theta_t(2)} \mathbf{f}_2$ is equal to $K \times (N + 1)$ and $\nabla_{\theta_t(2)} \mathbf{f}_2$ can be expressed as

$$\nabla_{\theta_t(2)} \mathbf{f}_2 = (\nabla_{\sigma_2} \mathbf{f}_2) \cdot [\mathbf{f}_1; 1]^T \quad (57)$$

where $\nabla_{\sigma_2} \mathbf{f}_2 = \frac{\partial \mathbf{f}_2[\mathbf{f}_1, \theta_t(2)]}{\partial \sigma_2}$.

Substituting $(\nabla_{\sigma_2} \mathbf{f}_2) \cdot [\mathbf{f}_1; 1]^T$ for $\nabla_{\theta_t(2)} \mathbf{f}_2$ in Equation 56, we can derive

$$\mathbf{f}_2[\mathbf{f}_1, \theta_{t+1}(2)] \approx \mathbf{f}_2[\mathbf{f}_1, \theta_t(2)] + (\nabla_{\sigma_2} \mathbf{f}_2) \cdot [\mathbf{f}_1; 1]^T \cdot \theta_{t+1}(2) - (\nabla_{\sigma_2} \mathbf{f}_2) \cdot [\mathbf{f}_1; 1]^T \cdot \theta_t(2) \quad (58)$$

If we only consider a single neuron, such as f_{2k} , in \mathbf{f}_2 , $\theta_{t+1}(2k) = [\beta_{1k}; \dots; \beta_{Nk}; b_{2k}]$ and $\theta_t(2k) = [\beta_{1k}^*; \dots; \beta_{Nk}^*; b_{2k}^*]$, thus $[\mathbf{f}_1; 1]^T \cdot \theta_{t+1}(2k) = \sum_{n=1}^N \beta_{nk} \cdot f_{1n} + b_{2k}$. As a result, Equation 58 with respect to a single neuron can be expressed as

$$\begin{aligned} f_{2k}[\mathbf{f}_1, \theta_{t+1}(2k)] &\approx (\nabla_{\sigma_2} f_{2k}) \cdot \left[\sum_{n=1}^N \beta_{nk} \cdot f_{1n} + b_{2k} \right] \text{ (First order approximation)} \\ &+ f_{2k}[\mathbf{f}_1, \theta_t(2k)] - (\nabla_{\sigma_2} f_{2k}) \cdot \left[\sum_{n=1}^N \beta_{nk}^* \cdot f_{1n} + b_{2k}^* \right] \text{ (Error)} \end{aligned} \quad (59)$$

Equation 59 indicates that $f_{2k}[\mathbf{f}_1, \theta_{t+1}(2k)]$ can be expressed as its first order approximation with an error component based on the activation in the previous iteration, i.e., $f_{2k}[\mathbf{f}_1, \theta_t(2k)]$. Since $\nabla_{\sigma_2} f_{2k} = \frac{\partial f_{2k}[\mathbf{f}_1, \theta_t(2)]}{\partial \sigma_2}$ is only related to \mathbf{f}_1 and the parameters in the t th training iteration, i.e., $\theta_t(2)$, it can be regarded as a constant. Also note the error component not contains any parameters in the $(t + 1)$ th training iteration. In summary, $f_{2k}(\mathbf{f}_1, \theta_{t+1}(2k))$ can be reformulated as

$$f_{2k}(\mathbf{f}_1, \theta_{t+1}(2k)) \approx C_1 \cdot \left[\sum_{n=1}^N \beta_{nk} \cdot f_{1n} + b_{2k} \right] + C_2 \quad (60)$$

where $C_1 = \nabla_{\sigma_2} f_{2k}$ and $C_2 = f_{2k}(\mathbf{f}_1, \theta_t(2k)) - (\nabla_{\sigma_2} f_{2k}) \cdot [\sum_{n=1}^N \beta_{nk}^* \cdot f_{1n} + b_{2k}^*]$. Similarly, the activations in the first hidden layer, i.e., \mathbf{f}_1 , also can be formulated as the first order approximation in the context of the gradient descent learning algorithm.

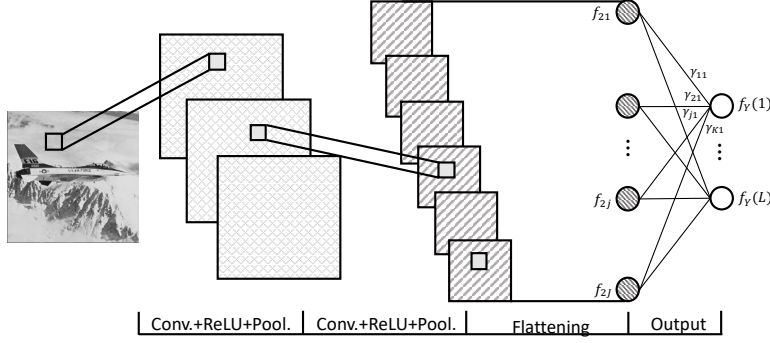


Figure 20: The above CNN has two convolutional layers, i.e., \mathbf{f}_1 , \mathbf{f}_2 , and one softmax output layer, i.e., \mathbf{f}_Y . We flatten the output of \mathbf{f}_2 as the input of \mathbf{f}_Y .

H THE PROPOSED PROBABILISTIC REPRESENTATION HOLDS FOR CONVOLUTIONAL LAYERS

This section proves that the proposed probabilistic representation is valid for convolutional layers. In the above CNN, the input \mathbf{x} has Q channels $\mathbf{x} = \{\mathbf{x}^q\}_{q=1}^Q$ (e.g., $Q = 3$ if \mathbf{x} are color images), \mathbf{f}_1 has N convolutional channels $\mathbf{f}_1 = \{\mathbf{f}_1^n\}_{n=1}^N$, \mathbf{f}_2 has K convolutional channels $\mathbf{f}_2 = \{\mathbf{f}_2^k\}_{k=1}^K$. The output layer \mathbf{f}_Y is the softmax with L nodes, thus its distribution can be formulated as

$$P(\mathbf{F}_Y) = \{P(f_{yl}) = \frac{1}{Z_{\mathbf{F}_Y}} \exp(f_{yl})\}_{l=1}^L \quad (61)$$

where $Z_{\mathbf{F}_Y} = \sum_{l=1}^L \exp(f_{yl})$ denotes the partition function. Since $\{f_{yl} = \sum_{j=1}^J \gamma_{jl} \cdot f_{2j} + b_{yl}\}_{l=1}^L$, we can derive

$$P(\mathbf{F}_Y) = \{P(f_{yl}) = \frac{1}{Z_{\mathbf{F}_Y}} \exp(\sum_{j=1}^J \gamma_{jl} \cdot f_{2j} + b_{yl})\}_{l=1}^L \quad (62)$$

where $\{f_{2j}\}_{j=1}^J$ is the flattened output of \mathbf{f}_2 , γ_{jl} is the weight of the edge between f_{2j} and f_{yl} , and b_{yl} denotes the bias. Since \mathbf{f}_2 is a convolutional layer with K convolutional channels, it can be reformulated as $\{f_{2j}\}_{j=1}^J = \{\mathbf{f}_2^k\}_{k=1}^K$. As a result, $P(\mathbf{F}_Y)$ can be reformulated as

$$P(\mathbf{F}_Y) = \{P(f_{yl}) = \frac{1}{Z_{\mathbf{F}_Y}} \exp(\sum_{k=1}^K \gamma_{kl}^T \cdot \mathbf{f}_2^k + b_{yl})\}_{l=1}^L \quad (63)$$

where γ_{kl} is a subset of all the parameters $\{\gamma_{jl}\}_{j=1}^J$ such that $\sum_{k=1}^K \gamma_{kl}^T \cdot \mathbf{f}_2^k = \sum_{j=1}^J \gamma_{jl} \cdot f_{2j}$, and γ_{kl}^T is the transpose of γ_{kl} . Therefore, $P(\mathbf{F}_Y)$ can be reformulated as

$$P(\mathbf{F}_Y) = \{P(f_{yl}) = \frac{1}{Z'_{\mathbf{F}_Y}} \prod_{k=1}^K \exp(\gamma_{kl}^T \cdot \mathbf{f}_2^k)\}_{l=1}^L \quad (64)$$

where $Z'_{\mathbf{F}_Y} = Z_{\mathbf{F}_Y} / [\exp(b_{yl})]$. Recall the element-wise matrix power, e.g., $\exp(\mathbf{a})^b$, where $\mathbf{a} = [a_1; a_2; a_3]$ and $\mathbf{b} = [b_1; b_2; b_3]$, we can derive that $\exp(\mathbf{a}) = [\exp(a_1); \exp(a_2); \exp(a_3)]$ and $\exp(\mathbf{a})^b = [\exp(a_1)^{b_1}; \exp(a_2)^{b_2}; \exp(a_3)^{b_3}] = [\exp(a_1 b_1); \exp(a_2 b_2); \exp(a_3 b_3)]$. As a result, $\exp(\mathbf{a}\mathbf{b}) = \exp(a_1 b_1 + a_2 b_2 + a_3 b_3) = \prod_{|a|} \exp(\mathbf{a})^b$, where $|a|$ is the element numbers of \mathbf{a} .

Based on the element-wise matrix power, we can reformulate $P(\mathbf{F}_Y)$ as

$$P(\mathbf{F}_Y) = \{P(f_{yl}) \approx \frac{1}{Z'_{\mathbf{F}_Y}} \prod_{k=1}^K \exp(\gamma_{kl}^T \cdot \mathbf{f}_2^k) = \frac{1}{Z'_{\mathbf{F}_Y}} \prod_{k=1}^K \prod_{|\mathbf{f}_2^k|} [\exp(\mathbf{f}_2^k)]^{\gamma_{kl}}\}_{l=1}^L \quad (65)$$

Moreover, we introduce a new partition function $Z_{F_2^k} = \sum_{f_2^k} \exp(\mathbf{f}_2^k)$ to guarantee that $\frac{1}{Z_{F_2^k}} \exp(\mathbf{f}_2^k)$ is a probability measure. As a result, $P(\mathbf{F}_Y)$ can be reformulated as

$$P(\mathbf{F}_Y) = \{P(f_{yl}) = \frac{1}{Z''_{F_Y}} \prod_{k=1}^K \prod_{|f_2^k|} [\frac{1}{Z_{F_2^k}} \exp(\mathbf{f}_2^k)]^{\gamma_{kl}}\}_{l=1}^L \quad (66)$$

where $Z''_{F_Y} = Z'_{F_Y} / \prod_{k=1}^K [Z_{F_2^k}]^{\gamma_{kl} \cdot |f_2^k|}$. Overall, $P(\mathbf{F}_Y)$ can be reformulated as a PoE model, in which each expert is defined by every convolutional channel in \mathbf{f}_2

$$P(\mathbf{f}_2^k) = \frac{1}{Z_{F_2^k}} \exp(\mathbf{f}_2^k) \quad (67)$$

Since \mathbf{f}_2 has K convolutional channels, i.e., $\mathbf{f}_2 = \{\mathbf{f}_2^k\}_{k=1}^K$, and each channel can be formulated as the summation of all the convolutional channels in \mathbf{f}_1 , i.e.,

$$\mathbf{f}_2 = \{\mathbf{f}_2^k = \sum_{n=1}^N \sigma_1(\mathbf{S}_2^{(k,n)} \circ \mathbf{f}_1^n) + b_1^k \cdot \mathbf{1}\}_{n=1}^N \quad (68)$$

where $\mathbf{S}_2^{(k,n)} \circ \mathbf{f}_1^n$ is the output of the k th convolutional filter applying into the n th channel of \mathbf{f}_1 , b_1^k is the bias, and $\sigma_1(\cdot)$ is the activation function.

If we regard all the convolutional filters for \mathbf{f}_2^k as a clique potential function, i.e. $\mathbf{f}_2^k = \varphi^k(\mathbf{f}_1)$, for modeling a signal structure of the input \mathbf{f}_1 , we can formulate $P(\mathbf{F}_2)$ as a specific Gibbs distribution, i.e., the Markov Random Fields (MRFs), in the context of the probabilistic graphical modeling.

$$P(\mathbf{F}_2) = \frac{1}{Z_{F_2}} \prod_{k=1}^K \exp(\mathbf{f}_2^k) = \frac{1}{Z_{F_2}} \exp(\sum_{k=1}^K \mathbf{f}_2^k) \quad (69)$$

where $Z_{F_2} = \sum_{f_2} \exp(\sum_{k=1}^K \mathbf{f}_2^k)$ is the partition function for $P(\mathbf{F}_2)$.

Based on Equation 67 and Equation 68, the distribution of an arbitrary channel in \mathbf{f}_2 , e.g, $P(\mathbf{f}_2^k)$, can be reformulated as

$$P(\mathbf{f}_2^k) = \frac{1}{Z'_{F_2^k}} \exp[\sum_{n=1}^N \sigma_1(\mathbf{S}_2^{(k,n)} \circ \mathbf{f}_1^n)] \quad (70)$$

where $Z'_{F_2^k} = Z_{F_2^k} / \exp(b_1^k)$.

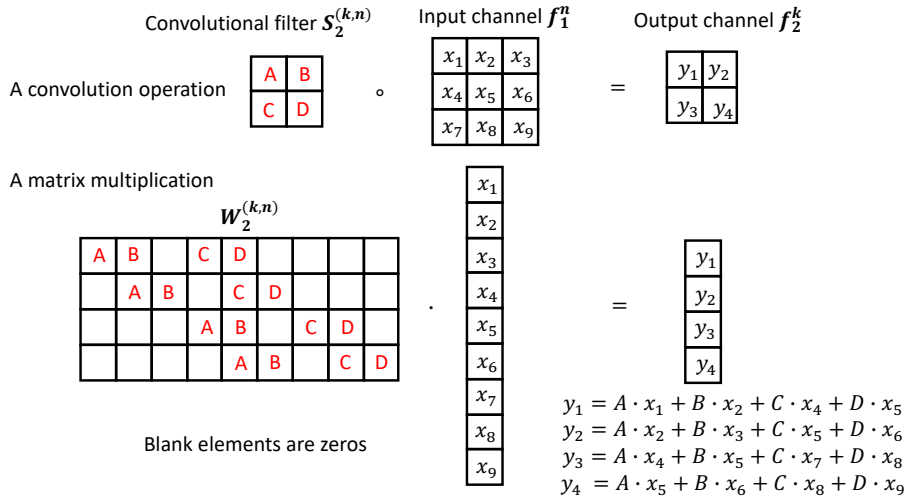


Figure 21: The equivalence between a convolutional operation and a full connected layer.

Based on the equivalence between the gradient descent learning and the first order approximation (Appendix G), we can derive

$$P(\mathbf{f}_2^k) \approx \frac{1}{Z'_{\mathbf{F}_2^k}} \exp\left(\sum_{n=1}^N \mathbf{S}_2^{(k,n)} \circ \mathbf{f}_1^n\right) \quad (71)$$

Based on the equivalence between a convolutional channel and a fully connected layer (Appendix E), we can regard a convolutional operation as a matrix multiplication, which can be visualized in Figure 21. As a result, we have

$$P(\mathbf{f}_2^k) \approx \frac{1}{Z'_{\mathbf{F}_2^k}} \exp\left(\sum_{n=1}^N \mathbf{W}_2^{(k,n)} \cdot \mathbf{f}_1^n\right) = \frac{1}{Z'_{\mathbf{F}_2^k}} \prod_{n=1}^N \exp(\mathbf{W}_2^{(k,n)} \cdot \mathbf{f}_1^n) \quad (72)$$

where $\mathbf{W}_2^{(k,n)}$ is a matrix corresponding to the convolution filter $\mathbf{S}_2^{(k,n)}$.

Based on the element-wise matrix power, $P(\mathbf{f}_2^k)$ can be reformulated as

$$P(\mathbf{f}_2^k) \approx \frac{1}{Z'_{\mathbf{F}_2^k}} \prod_{n=1}^N \exp(\mathbf{W}_2^{(k,n)} \cdot \mathbf{f}_1^n) = \frac{1}{Z'_{\mathbf{F}_2^k}} \prod_{n=1}^N \exp(\mathbf{f}_1^n) \mathbf{W}_2^{(k,n)} \quad (73)$$

Similarly, we introduce a new partition function $Z_{\mathbf{F}_1^n} = \sum_{\mathbf{f}_1^n} \exp(\mathbf{f}_1^n)$ to guarantee that $\frac{1}{Z_{\mathbf{F}_1^n}} \exp(\mathbf{f}_1^n)$ is a probability measure. As a result, $P(\mathbf{f}_2^k)$ can be reformulated as

$$P(\mathbf{f}_2^k) = P(f_{yl}) \approx \frac{1}{Z''_{\mathbf{F}_2^k}} \prod_{n=1}^N \left[\frac{1}{Z_{\mathbf{F}_1^n}} \exp(\mathbf{f}_1^n) \right] \mathbf{W}_2^{(k,n)} \quad (74)$$

where $Z''_{\mathbf{F}_2^k} = Z'_{\mathbf{F}_2^k} / \prod_{n=1}^N [Z_{\mathbf{F}_1^n}] \mathbf{W}_2^{(k,n)}$. Therefore, $P(\mathbf{f}_2^k)$ can be reformulated as a PoE model, in which all the experts are defined by the convolutional channels in the previous convolutional layer.

Similar as Equation 69, the distribution of \mathbf{f}_1 also can be expressed as a MRF model

$$P(\mathbf{F}_1) = \frac{1}{Z_{\mathbf{F}_1}} \prod_{n=1}^N \exp(\mathbf{f}_1^n) = \frac{1}{Z_{\mathbf{F}_1}} \exp\left(\sum_{n=1}^N \mathbf{f}_1^n\right) \quad (75)$$

where $Z_{\mathbf{F}_1} = \sum_{\mathbf{f}_1} \exp(\sum_{n=1}^N \mathbf{f}_1^n)$ is the partition function for $P(\mathbf{F}_1)$.

Overall, the distribution of a convolutional layer can be formulated as a MRF model.

I THE APPROACH TO GENERATE THE SYNTHETIC DATASET BASED ON NIST AND THE ARCHITECTURE OF THE CNN

In this section, we present a novel approach to generate a synthetic dataset obeying the Gaussian distribution based on the NIST¹ dataset of handwritten digits by class. The synthetic dataset has similar characteristics as the benchmark MNIST dataset. It consists of 20,000 32×32 grayscale images in 10 classes (digits from 0 to 9), and each class has 1,000 training and 1,000 testing images. More specifically, the approach to generate the synthetic dataset is summarized in Algorithm 1, and Figure 22 visualize the spatial relation of $\hat{\mathbf{x}} = \{\hat{\mathbf{x}}_{\text{outside}}, \hat{\mathbf{x}}_{\text{outside-boundary}}, \hat{\mathbf{x}}_{\text{inside-boundary}}, \hat{\mathbf{x}}_{\text{inside}}\}$.

¹<https://www.nist.gov/srd/nist-special-database-19>

Algorithm 1 The algorithm for generating the synthetic dataset**input:** NIST dataset of handwritten digits by class

- 1: **repeat**
- 2: binarizing an image of NIST to obtain z
- 3: extracting the central part of z to obtain z_c with dimension 64×64
- 4: downsampling z_c to obtain z_{cd} with dimension 32×32
- 5: extracting the edge of z_{cd} to obtain the mask image m_{cd}
- 6: decomposing m_{cd} into four parts, i.e., m_{outside} , $m_{\text{outside-boundary}}$, $m_{\text{inside-boundary}}$, and m_{inside} .
- 7: sampling $\mathcal{N}(0, 1024)$ to derive a random vector $x \in \mathbb{R}^{1024 \times 1}$
- 8: sorting x in the descending order to derive \hat{x}
- 9: decomposing \hat{x} into four parts, i.e., $\hat{x} = \{\hat{x}_{\text{outside}}, \hat{x}_{\text{outside-boundary}}, \hat{x}_{\text{inside-boundary}}, \hat{x}_{\text{inside}}\}$
- 10: placing each pixel of $\hat{x} = \{\hat{x}_{\text{outside}}, \hat{x}_{\text{outside-boundary}}, \hat{x}_{\text{inside-boundary}}, \hat{x}_{\text{inside}}\}$ into the above masks.
- 11: **until** (20,000 synthetic images are generated)

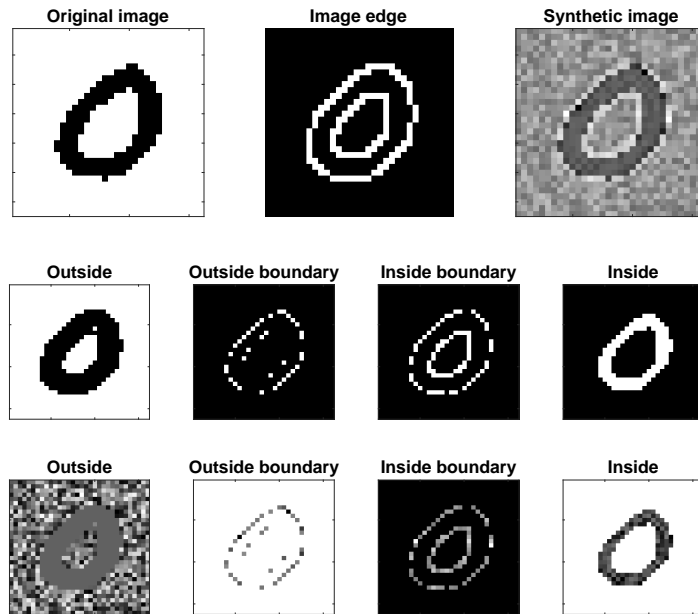
output: The synthetic dataset

Figure 22: The first row shows an original image, its edge, and the corresponding synthetic image based on the original one. The second row uses white pixels to show the spatial position of the four mask parts, i.e., m_{outside} , $m_{\text{outside-boundary}}$, $m_{\text{inside-boundary}}$, and m_{inside} . The third row shows the synthetic image corresponding to each mask part, i.e., $\hat{x} = \{\hat{x}_{\text{outside}}, \hat{x}_{\text{outside-boundary}}, \hat{x}_{\text{inside-boundary}}, \hat{x}_{\text{inside}}\}$.

The architecture of the CNN for classifying the synthetic dataset is summarized in Table 8.

Table 8: The architectures of the CNN for synthetic image classification

R.V.	Layer	Description	CNN
X	x	Input	$32 \times 32 \times 1$
F_1	f_1	Conv (3×3) + ReLU	$30 \times 30 \times 20$
	f_2	Maxpool	$15 \times 15 \times 20$
F_2	f_3	Conv (5×5) + ReLU	$11 \times 11 \times 60$
	f_4	Maxpool	$5 \times 5 \times 60$
F_Y	f_5	Output(softmax)	$1 \times 1 \times 10$

R.V. is the random variable of the hidden layer(s).




Figure 23: The training procedure of a DNN = $\{\mathbf{x}; \mathbf{f}_1; \dots; \mathbf{f}_Y\}$. The red arrow indicates the traditional training algorithm, e.g. backpropagation, which derives an output of the DNN given training dataset $\{\mathbf{x}^j\}_{j=1}^J$ in the forward direction and calculate the gradient of parameters with respect to the training labels \mathbf{Y} for updating the parameters in the backward direction.

J A REGULARIZATION ALGORITHM BASED ON THE PROBABILISTIC REPRESENTATION

J.1 THE PROPOSED REGULARIZATION ALGORITHM

Based on the probabilistic representation, the first hidden layer formulates a MRF model as a prior distribution of the training dataset and other hidden layers recursively serve as prior distributions for their respective next layer in a CNN. However, the backpropagation calculates the gradient in the backward direction, i.e., updating the gradient of a hidden layer corresponding to a prior distribution must go over all hidden layers behind it. In addition, the degradation problem could make the learned prior distribution not precise enough to regularize DNNs and results in overfitting, especially in very deep neural networks He et al. (2016).

In this section, we propose a new regularization learning algorithm for DNNs, which includes two steps: (i) we learn the prior distributions of DNNs directly from the training dataset $\{\mathbf{x}^j\}_{j=1}^J$ as the blue arrow shown in Figure 23 and (ii) we use the backpropagation to train the initialized DNNs by the learned prior distributions as the red arrow shown in Figure 23.

Since directly learning prior distribution from $\{\mathbf{x}^j\}_{j=1}^J$ is less complicated than the backpropagation learning algorithm, we can relieve the effects of degradation problem and derive prior distributions more precisely, thereby achieving the better generalization performance of DNNs.

An important question need to answer is how many hidden layers of DNNs need to be pre-trained in the first step. Since most prior information is stored in the training dataset $\{\mathbf{x}^j\}_{j=1}^J$, the hidden layer closest to the input of DNNs, i.e., the first hidden layer \mathbf{f}_1 , should learn the most prior information from $\{\mathbf{x}^j\}_{j=1}^J$ as long as \mathbf{f}_1 has powerful representation ability through including sufficient hidden units, e.g., convolutional filters. Therefore, the proposed regularization method only pre-train the first hidden layer \mathbf{f}_1 in order to decrease the computation complexity. Pre-training more hidden layers could obtain better generalization performance, which can be a direction for future research.

We choose the Field of Experts (FoE) model to directly learn a prior distribution of $\{\mathbf{x}^j\}_{j=1}^J$ because of two reasons. Above all, FoE achieves great performance on learning prior knowledge of given datasets on various applications, e.g., image denoising Roth & Black (2005), image restoration Roth & Black (2008), and image inpainting Schmidt et al. (2011). More importantly, FoE can be regarded as a Convolutional Neural Networks (CNN) with a single convolutional layer, hence we can easily use the learned FoE model to initialize the first hidden layer of CNNs as long as they have the same architecture. A FoE model can be formulated as

$$p_{\text{FoE}}(\mathbf{x}) = \frac{1}{Z_{\text{FoE}}} \prod_{k=1}^K \{f_k^{\text{expert}}[f_k(\mathbf{x})]\}, \quad (76)$$

where $f_k(\cdot)$ is a linear filter, f_k^{expert} is an expert function, and $Z_{\text{FoE}} = \int_{\mathbf{x}} \prod_{k=1}^K \{f_k^{\text{expert}}[f_k(\mathbf{x})]\} d\mathbf{x}$ is the partition function. We use the contrastive divergence learning method to train $p_{\text{FoE}}(\mathbf{x})$ under the Kullback-Leibler divergence (KLD) criterion (Carreira-Perpinan & Hinton, 2005). More detail of the training procedure is in (Roth & Black, 2005; 2008). After deriving a well-trained FoE model, we use the learned linear filters $\{f_k\}_{k=1}^K$ to initialize the first convolutional layer of CNNs. Finally, we use the backpropagation algorithm to train the initialized CNN for a specific application, such as image recognition. We summarize the proposed regularization learning algorithm in Algorithm 2.

Algorithm 2 Regularization learning algorithm

input: $\mathcal{D} = \{(\mathbf{x}_n, \mathbf{y}_n)\}_{n=1}^N$, DNN = $\{\mathbf{x}; \mathbf{f}_1; \dots; \mathbf{f}_I; \mathbf{f}_Y\}$, and $p_{\text{FoE}}(\mathbf{x})$

- 1: **setup**
- 2: initialize iteration $k = 0$, and training epochs K
- 3: specify the architecture of the DNN and $p_{\text{FoE}}(\mathbf{x})$
- 4: **regularization learning**
- 5: train $p_{\text{FoE}}(\mathbf{x})$ given $\{\mathbf{x}_n\}_{n=1}^N$
- 6: initialize \mathbf{f}_1 by the well-trained $p_{\text{FoE}}(\mathbf{x})$
- 7: **repeat**
- 8: train the DNN given \mathcal{D}
- 9: **until** ($k = K$)

output: the well-trained DNN for image recognition

J.2 SIMULATIONS

In this section, we validate the proposed regularization learning algorithm for image recognition application on the CIFAR-10 benchmark dataset, which includes 70,000 32×32 color images in 10 classes, and each class has 6,000 training images and 1,000 testing images. Specifically, we randomly choose 2,000 training images of each class as a new training dataset and use the same testing dataset. Smaller training dataset means overfitting being more likely to occur, which can help us validate the proposed regularization learning method more convincingly. In addition, we covert color images to grayscale to simplify computation.

In order to directly learn regularization from $\{\mathbf{x}^j\}_{j=1}^{20000}$, we design a FoE model with 12 Gaussian Scale Mixture (GSM) experts, thus it can be expressed as

$$p_{\text{FoE}}(\mathbf{x}) = \frac{1}{Z_{\text{FoE}}} \prod_{k=1}^{12} \{f_k^{\text{GSM}}[f_k(\mathbf{x})]\}, \quad (77)$$

where $f_k^{\text{GSM}}[f_k(\mathbf{x})] = \sum_{i=1}^{11} \pi_{ki} \cdot \mathcal{N}(f_k(\mathbf{x}); 0, \frac{\sigma_b^2}{\delta(i)})$, σ_b^2 is a fixed base variance, $\delta = \{\delta(1), \dots, \delta(11)\}$ is a range of scales, $f_k(\cdot)$ is a 3×3 convolutional filter, and π_{ki} denotes the weight of each Gaussian distribution. The parameters of the FoE model is $\theta = \{\mathbf{f}, \boldsymbol{\pi}\}$, where $\mathbf{f} = [f_1, \dots, f_{12}]$ and $\boldsymbol{\pi} = \{\pi_{ki}\}_{k=1, i=1}^{k=12, i=11}$.




Figure 24: Some grayscale images derived from the benchmark CIFAR-10 dataset.

Table 9: The architectures of CNN for CIFAR-10 classification

Layer	Description	Output
x	Input	$32 \times 32 \times 1$
f_1	Conv (3×3)	$30 \times 30 \times 12$
f_2	ReLU	$30 \times 30 \times 12$
f_3	Conv (5×5)	$26 \times 26 \times 32$
f_4	ReLU + Maxpool	$13 \times 13 \times 32$
f_5	Conv (5×5)	$9 \times 9 \times 128$
f_6	ReLU + Maxpool	$4 \times 4 \times 128$
f_7	Fully connected	$1 \times 1 \times 512$
f_8	ReLU	$1 \times 1 \times 512$
f_9	Fully connected	$1 \times 1 \times 32$
f_{10}	ReLU	$1 \times 1 \times 32$
f_Y	Output(softmax)	$1 \times 1 \times 10$

The architecture of the CNN for image recognition is summarized in Table 9. It is noteworthy that the number of convolutional filters in f_1 and the filter dimension are the same as the above FoE model, thus we can use the well-trained FoE to initialize f_1 as a regularization for the CNN classifying the CIFAR-10 dataset based on Algorithm 2.

In order to make a comprehensive comparison, we choose two commonly used regularizations, i.e., the dropout (drop rate is 0.3) and the L_2 norm, as references. We use the testing error to exam the generalization performances of different regularizations. Figure 25 visualizes that the proposed regularization method (abbr. Bayes Prior (BP)) successfully decreases the testing error. In particular, it outperforms dropout (32.26%) and achieve similar performance as L_2 norm (29.58%). Therefore, this experiment validates the effectiveness of the proposed regularization learning method.

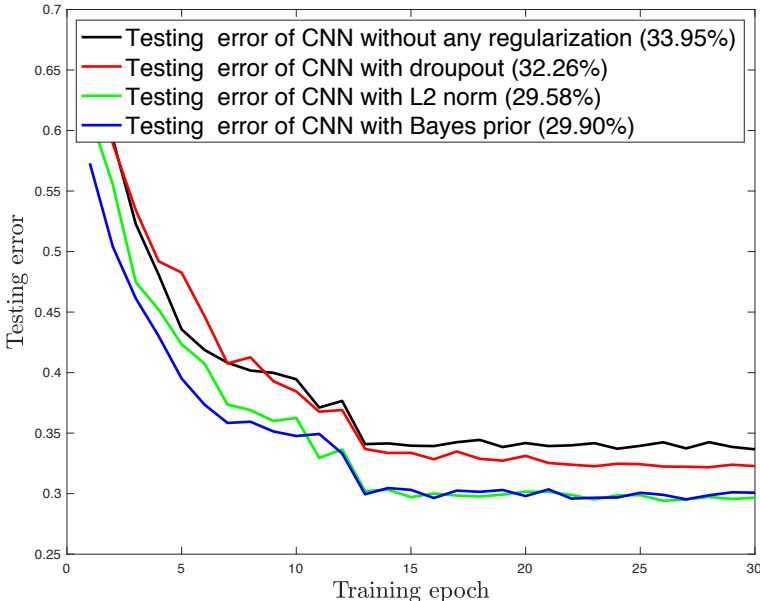


Figure 25: The average testing error of the CNN on the last 5 training epochs based on various regularizations. Bayes prior indicates the proposed regularization and its testing error is 29.90%.

Moreover, the proposed regularization method can be easily combined with other regularizations methods to improve the generalization performance further. We combine two arbitrary aforementioned regularization methods together and examine their generalization performances under the same experimental conditions as above. Figure 26 shows that the combination of two regularizations together achieves better generalization performance than any single one. In particular, the combination of BP and L2 norm achieves the lowest testing error (28.60%). Also note that if we combine three regularizations together, we achieve the best generalization performance (27.69%). This experiment further validate the proposed regularization learning algorithm.

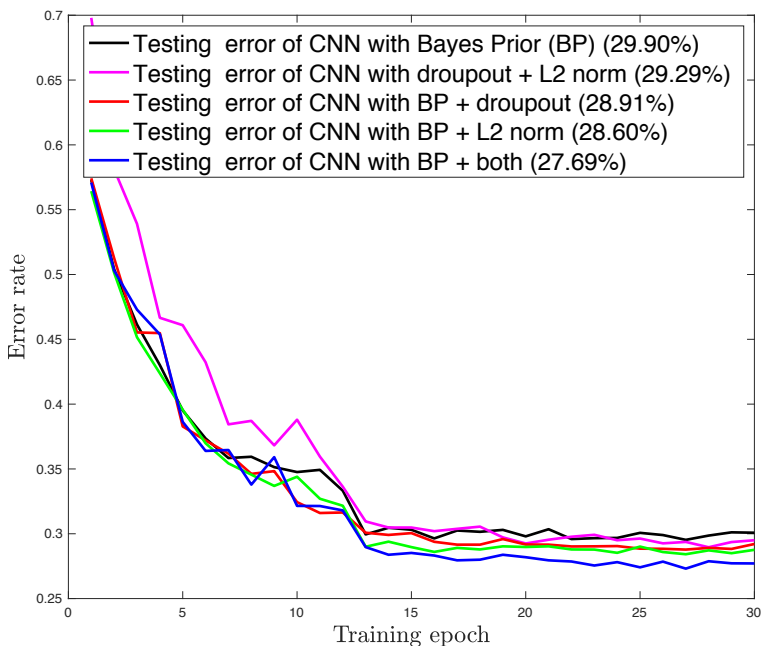


Figure 26: The average testing error of the CNN on the last 5 training epochs based on various combined regularizations. Bayes prior + both indicates that we combine all three regularizations together.