
Understanding Representation of Deep Equilibrium Models from Neural Collapse Perspective

Haixiang Sun
ShanghaiTech University
sunhx@shanghaitech.edu.cn

Ye Shi *
ShanghaiTech University
shiye@shanghaitech.edu.cn

Abstract

Deep Equilibrium Model (DEQ), which serves as a typical implicit neural network, emphasizes their memory efficiency and competitive performance compared to explicit neural networks. However, there has been relatively limited theoretical analysis on the representation of DEQ. In this paper, we utilize the Neural Collapse (\mathcal{NC}) as a tool to systematically analyze the representation of DEQ under both balanced and imbalanced conditions. \mathcal{NC} is an interesting phenomenon in the neural network training process that characterizes the geometry of class features and classifier weights. While extensively studied in traditional explicit neural networks, the \mathcal{NC} phenomenon has not received substantial attention in the context of implicit neural networks. We theoretically show that \mathcal{NC} exists in DEQ under balanced conditions. Moreover, in imbalanced settings, despite the presence of minority collapse, DEQ demonstrated advantages over explicit neural networks. These advantages include the convergence of extracted features to the vertices of a simplex equiangular tight frame and self-duality properties under mild conditions, highlighting DEQ's superiority in handling imbalanced datasets. Finally, we validate our theoretical analyses through experiments in both balanced and imbalanced scenarios.

1 Introduction

Recently, there has been significant research on *implicitly-defined layers* in neural networks [1, 2, 3, 6, 9, 20, 21, 49], where the output is implicitly mapped from the input under certain conditions. These layers embed interpretability and introduce inductive bias [26] into black-box neural networks, demonstrating superior performance compared to existing explicit layers.

Among these implicit networks, the Deep Equilibrium Model (DEQ) is a memory-efficient architecture that represents all hidden layers as the equilibrium point of a nonlinear fixed-point equation. Due to the absence of a closed-form solution in its forward process, DEQ can be viewed as having an infinite number of layers during iteration as long as the threshold is set low enough, enhancing its ability to fit input data. Consequently, its representational capacity is relatively stronger compared to a single-layer network structure. This phenomenon explains why DEQ has achieved state-of-the-art results in classification tasks compared to existing architectures like ResNet. For instance, it has been successfully applied to language tasks and image classification tasks, reaching state-of-the-art performance. Additionally, DEQ can be applied in various domains and integrated with numerous other models, including inverse problems [19], Neural ODEs [11], diffusion models [24, 43], Gaussian processes [17], and more.

However, recent research reveals a phenomenon called Neural Collapse (\mathcal{NC}) concerning the learned deep representations across datasets in image classification tasks [42]. Under the \mathcal{NC} regime, the last-

*Corresponding author.

layer feature of each sample in neural networks collapses to their within-class mean, and the classifier vector converges to a simplex Equiangular Tight Frame (ETF). Theoretical analyses [10, 40, 51] indicate that under the Unconstrained Features Mode (UFM) condition, specific features \mathbf{H}^0 can be isolated from the entire network, known as the layer-peeled model [14]. In this scenario, Neural Collapse (\mathcal{NC}) is observed under certain conditions, suggesting that \mathcal{NC} is agnostic to the backbone of feature extraction. Moreover, since \mathcal{NC} measures the degree of proximity between features of the same category, an imbalanced dataset can exert a more negative influence on the performance of \mathcal{NC} . For instance, classes with fewer samples may not separate well and could converge in the same direction, leading to what is known as *Minority Collapse* [14]. Thus, the \mathcal{NC} metric serves as a valuable indicator of a model's behavior in the context of imbalanced datasets.

The reasons behind the superior performance of DEQ still lack theoretical proof and comprehensive quantitative analysis. Additionally, to the best of our knowledge, no prior work has integrated DEQ with imbalanced scenarios. In our study, we integrate DEQ with layer-peeled models, add constraints with respect to weights \mathbf{W}_{DEQ} , and consider the results of fixed-point iteration as the output of DEQ. Therefore, we analyze the performance of \mathcal{NC} in DEQ by continuously deriving the lower bound of the loss function under certain constraints, allowing us to assess how \mathcal{NC} manifests in the training performance of the network. Similarly, we apply the same operations to explicit neural networks for comparison. Our results show that DEQ performs similarly to explicit neural networks under balanced settings. We further extend the dataset to imbalanced conditions and analyze the \mathcal{NC} performance in DEQ, explaining why DEQ tends to outperform explicit neural networks under mild conditions. We systematically analyze performance in terms of feature convergence, distance to the Simplex ETF, and the parallel relationship between extracted features and classifier weights. These analyses uncover the reasons behind the superior performance of DEQ compared to explicit neural networks during training. Additionally, the experimental results in both balanced and imbalanced scenarios validate our theoretical analyses.

Our main contributions are:

- We systematically analyzed the representation of DEQ from the \mathcal{NC} perspective and compared their performance with explicit neural networks. Our theoretical analysis shows that both DEQ and explicit neural networks exhibit the \mathcal{NC} phenomenon in balanced datasets.
- Under imbalanced settings, we theoretically proved the convergence of extracted features to the vertices of a simplex ETF and alignment with classifier weights under certain conditions, demonstrating DEQ's advantages over explicit neural networks under some mild conditions.
- Experimental results on Cifar-10 and Cifar-100 validated our theoretical findings for distinguishing the differences between DEQ and explicit neural networks.

2 Background and related works

We consider a classification task with K classes. Let n_k denote the number of training samples in each class k , and $N = \sum_{k=1}^K n_k$ represent the total number of training samples. A traditional neural network can be expressed as a mapping:

$$\psi(\mathbf{x}) = \mathbf{W}\phi(\mathbf{x}) + \mathbf{b}, \quad (1)$$

where $\phi(\mathbf{x}) : \mathbb{R}^{\text{in} \times N} \rightarrow \mathbb{R}^{D \times N}$ is the feature extraction, $\mathbf{W} \in \mathbb{R}^{K \times D}$ and $\mathbf{b} \in \mathbb{R}^K$ are the classifiers and bias in the last layer, respectively. For simplicity, we consider the bias-free case and omit the term \mathbf{b} . Besides, we will denote $\mathbf{H} = \phi(\mathbf{x})$ in later sections.

2.1 Deep Equilibrium Models

There have been numerous neural network architectures designed for various practical tasks from different perspectives [17, 32, 38, 39, 48]. DEQ, a typical implicit network [13, 52], incorporates unrolling methods [12, 41], which are devised for training arbitrarily deep networks by integrating all the network layers into one [3, 4, 5, 35, 36, 58].

Let $f_\theta(z, \mathbf{x})$ represent a DEQ layer with input \mathbf{x} parameterized by θ . When z^* reaches the equilibrium point, it satisfies:

$$g_\theta(z^*, \mathbf{x}) \triangleq f_\theta(z^*, \mathbf{x}) - z^* = 0. \quad (2)$$

The forward procedure mostly employs the Broyden solver [7] for iterative solving:

$$z_{t+1} = z_t - \mathbf{B}_t^{-1} g_\theta(z_t, \mathbf{x}), \quad (3)$$

where \mathbf{B}_t^{-1} refers to the approximation of inverse matrix $\nabla_z^{-1} g_\theta(z_t, \mathbf{x})$, as well as the same parameter θ shared across iterations. However, the solution can be quite unstable, and efforts have been made to enhance stability and robustness [34, 44, 55, 56]. Especially, regarding the computation of the inverse matrix, it can be expanded in the form of a Neumann series [18, 60]. Besides, accelerating and stabilizing the backward procedure is also an important issue in DEQ [15].

2.2 Neural Collapse \mathcal{NC}

The phenomenon of \mathcal{NC} was initially uncovered by [42], which is considered an intriguing regularity in neural network training with many elegant geometric properties [50, 61, 66]. When the model is at the terminal phase of training (TPT), or more precisely, achieves zero training error, the within-class means of features and the classifier vectors converge to the vertices of a simplex Equiangular Tight Frame (ETF) on a balanced dataset.

Definition 2.1. (Simplex Equiangular Tight Frame) A collection of points $\mathbf{s}_i \in \mathbb{R}^D, i = 1, 2, \dots, K$, is said to be a simplex equiangular tight frame if

$$\mathbf{S} = \alpha \sqrt{\frac{K}{K-1}} \mathbf{P} \left(\mathbf{I}_K - \frac{1}{K} \mathbf{1}_K \mathbf{1}_K^T \right), \quad (4)$$

where α is a non-zero scalar, $\mathbf{S} = [\mathbf{s}_1, \dots, \mathbf{s}_k] \in \mathbb{R}^{D \times K}$, $\mathbf{I}_K \in \mathbb{R}^{K \times K}$ is the identity matrix, $\mathbf{1}_K$ is the ones vector, and $\mathbf{P} \in \mathbb{R}^{D \times K} (D \geq K)$ is a partial orthogonal matrix such that $\mathbf{P}^T \mathbf{P} = \mathbf{I}_K$.

\mathcal{NC} incorporates the following four properties of the last-layer features and classifiers in deep learning training on balanced datasets:

$\mathcal{NC}1$: **Variability collapse:** The feature within-class converges to a unique vector, *i.e.*, for any sample i in the same class k , its feature $\mathbf{h}_{k,i}$ satisfies $\|\mathbf{h}_{k,i} - \bar{\mathbf{h}}_k\| \rightarrow 0, k \in [k]$, with the training procedure.

$\mathcal{NC}2$: **Convergence to simplex ETF:** The mean value \mathbf{h}^* of optimal features for each class collapses to the vertices of the simplex ETF.

$\mathcal{NC}3$: **Convergence to self-duality:** The class means and the classifier weights mutually converge: $\frac{\mathbf{W}^*}{\|\mathbf{W}^*\|} = \frac{\mathbf{H}^*}{\|\mathbf{H}^*\|}$.

$\mathcal{NC}4$: **Nearest Neighbor:** The classifier determines the class based on the Euclidean distances among the feature vector and the classifier weights.

2.3 Layer-peeled model under balanced and imbalanced conditions

Current studies often focus on the case where only the last-layer features and classifier are learnable without considering the layers in the backbone network under the assumption of Unconstrained Features Mode (UFM) [66], which can also be referred to as the Layer-peeled Model [14, 28]. First, we define the feasible set of parameters:

$$\mathcal{C} = \left\{ \mathbf{w}_k, \mathbf{h}_{k,i} \mid \frac{1}{K} \sum_{k=1}^K \|\mathbf{w}_k\|^2 \leq E_W, \frac{1}{K} \sum_{k=1}^K \frac{1}{n_k} \sum_{i=1}^{n_k} \|\mathbf{h}_{k,i}\|^2 \leq E_H \right\}. \quad (5)$$

Definition 2.2. (Layer-peeled Model) When \mathbf{H} and \mathbf{W} are the last layer classifier and weights respectively, then the optimization process of the neural network can be reformulated as:

$$\min_{\mathbf{W}, \mathbf{H}} \frac{1}{N} \sum_{k=1}^K \sum_{i=1}^{n_k} \mathcal{L}(\mathbf{W} \mathbf{h}_{k,i}, \mathbf{y}_k) \quad \text{s.t.} \quad \mathbf{w}_k, \mathbf{h}_{k,i} \in \mathcal{C}, \quad (6)$$

where E_H and E_W are two predefined values, N refers to the total number of samples.

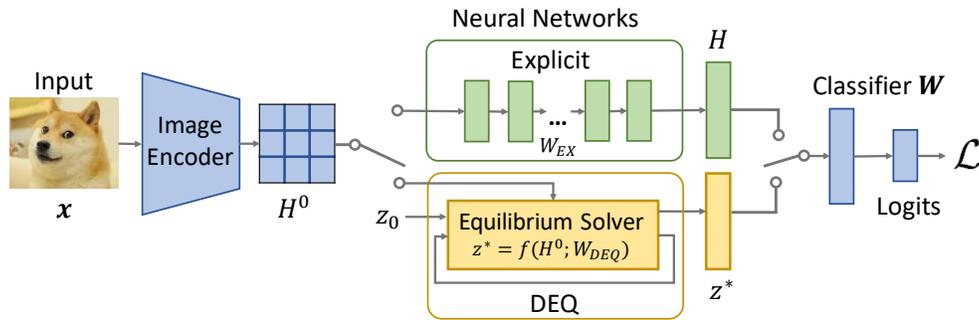


Figure 1: Illustration of feature extraction. After extracting feature maps H^0 , further features H or z^* can be obtained by passing through an explicit neural network or DEQ. The final step involves the classifier to obtain predicted logits. To ensure a fair comparison, we standardize the backbone network and its output H^0 across all conditions.

It should be noted that all the loss functions \mathcal{L} analyzed in our study are cross-entropy, as most current research focuses on this widely used deep learning classification loss function [23, 31]. And though the optimization program is nonconvex; however, it can generally be mathematically tractable for analysis. Besides, experiments with unregularized loss function and randomly initialized gradient descent typically converge to non-collapse global minimizers [51].

Under UFM, most \mathcal{NC} studies are based on 1-2 conventional layers of weights, However, there is also work [10, 51] that extends it to analyze M linear layers. Additionally, various studies have revealed additional characteristics of \mathcal{NC} , such as its impact on generalization [16, 25, 27, 61], its influence on feature learning [45], global optimality of the network [64, 66] and others. Therefore, \mathcal{NC} is a very efficient tool to analyze the performance of neural networks.

Imbalanced learning However, \mathcal{NC} will not occur under imbalanced settings generally. This phenomenon arises due to the imbalance in sample quantities, leading to challenges in adequately fitting features for certain classes. This is commonly referred to as *minority collapse* [8, 14]. As the degree of imbalance increases, it is expected that classifiers for minority classes converge. When Minority Collapse occurs, the neural network predicts equal probabilities for all minority classes, regardless of the input.

To enhance learning performance in imbalanced scenarios [62] and mitigate the effects of minority collapse, several methods have been proposed. [14] introduced convex relaxation, modifying a loss function [57], and incorporating a regularization term [37]. The reweighted approach is also widely applied, with some studies measuring it based on sample quantities [47, 59]. Additionally, adaptive techniques such as AutoBalance [33] have been introduced, which incorporates a bilevel optimization framework, along with logit balance [46, 54, 63, 65].

3 Comparison under balanced setting

In this section, we first analyze the \mathcal{NC} phenomenon in DEQ under balanced settings. As illustrated in Figure 1, after completing the initial feature extraction, we further examine the feature H obtained respectively by explicit neural networks and DEQ to reveal the \mathcal{NC} phenomenon.

3.1 \mathcal{NC} in Explicit Neural Networks

Building upon (6), we analyze \mathcal{NC} in explicit neural networks by considering the following constrained optimization problem during training:

$$\begin{aligned} \min_{\mathbf{W}, \mathbf{W}_{EX}, H^0} \quad & \frac{1}{N} \sum_{k=1}^K \sum_{i=1}^n \mathcal{L}(\mathbf{W} \mathbf{W}_{EX} h_{k,i}^0, \mathbf{y}_k) \\ \text{s.t.} \quad & \|\mathbf{W}_{EX}\|_F \leq E_H; \mathbf{w}_k, \mathbf{h}_{k,i} \in \mathcal{C}, \end{aligned} \quad (7)$$

where each n_k is set to n under the balanced setting, \mathbf{W}_{EX} represents the subsequent network weights. For ease of comparison with DEQ, we assume that the final feature is represented as $\mathbf{H} = \mathbf{W}_{EX}\mathbf{H}^0$. Traditional neural network structures are nonconvex, making them challenging to analyze due to their highly interactive nature. Employing the layer-peeled model alleviates the difficulty of \mathcal{NC} analysis.

3.2 \mathcal{NC} in Deep Equilibrium models

Building upon recent investigations into the \mathcal{NC} phenomenon, we embrace the layer-peeled model, where the last-layer features $\mathbf{h} = \phi(\mathbf{x})$ (equilibrium points in DEQ \mathbf{z}^*) as unconstrained optimization variables. Accordingly, we add the following constraints to enforce \mathcal{NC} in DEQ:

$$\mathcal{C}_{DEQ} \triangleq \{\mathbf{z}^*, \mathbf{W}_{DEQ} \mid \mathbf{z}^* = f(\mathbf{H}^0; \mathbf{W}_{DEQ}), \|\mathbf{W}_{DEQ}\|_F \leq E_H\}. \quad (8)$$

Compared to explicit layers, the active parameter in Deep Equilibrium models is \mathbf{W}_{DEQ} , hence imposing restrictions on it to align with the same feasible space. Then the formulation of DEQ with \mathcal{NC} becomes:

$$\begin{aligned} \min_{\mathbf{W}, \mathbf{W}_{DEQ}, \mathbf{z}^*, \mathbf{H}^0} \quad & \frac{1}{N} \sum_{k=1}^K \sum_{i=1}^{n_k} \mathcal{L}(\mathbf{W}\mathbf{z}^*, \mathbf{y}_k) \\ \text{s.t.} \quad & \mathbf{w}_k, \mathbf{h}_{k,i} \in \mathcal{C}; \mathbf{z}^*, \mathbf{W}_{DEQ} \in \mathcal{C}_{DEQ}. \end{aligned} \quad (9)$$

No matter whether under DEQ or explicit neural networks, these constraints must be imposed. This is because when these constraints are satisfied and the loss function reaches its lower bound, the \mathcal{NC} phenomenon is guaranteed. In our theoretical analysis, we assume that the DEQ is linear, that is, $\mathbf{z}^* = \text{fixed-point}(f_\theta(\mathbf{x}), \mathbf{z}) = \sum_{i=0}^{\infty} \mathbf{W}_{DEQ}^i \mathbf{x}$. Detailed analysis incorporating these constraints is provided in Appendix B.

The following theorem elucidates the specific scenarios in which the \mathcal{NC} phenomenon occurs. For a fair comparison, we assume that the extracted features \mathbf{H}^0 of the image encoder are the same in the derivation.

Theorem 3.1. (Feature collapse of explicit fully connected layers and implicit deep equilibrium models under balanced setting) Suppose (7) and (9) reaches its minimal, then

$\mathcal{NC}1$: For $\forall k = 1, 2, \dots, K$ and $\forall i = 1, 2, \dots, n$:

$$\mathbf{W}_{EX}\mathbf{h}_{k,i}^0 = \mathbf{W}_{EX}\mathbf{h}_k^0,$$

where $\mathbf{h}_k^0 = \sum_{i \in \tau(k)} \mathbf{h}_{k,i}^0$. Similarly, if the model is DEQ, then

$$f(\mathbf{h}_{k,i}^0; \mathbf{W}_{DEQ}) = f(\mathbf{h}_k^0; \mathbf{W}_{DEQ}).$$

$\mathcal{NC}2$: The classifier aligns to the Simplex ETF, regardless of whether explicit neural network and DEQ are applied:

$$\begin{aligned} \mathbf{W}\mathbf{W}^T &= \sqrt{E_W/E_H} \mathbf{W}\mathbf{W}_{EX}\mathbf{H}^0 \\ &= \sqrt{E_W/E_H} \mathbf{W}f(\mathbf{H}^0; \mathbf{W}_{DEQ}) \\ &= \frac{KE_W}{K-1} \left(\mathbf{1}_K - \frac{1}{K} \mathbf{1}_K \mathbf{1}_K^T \right). \end{aligned}$$

$\mathcal{NC}3$: For $\forall k = 1, 2, \dots, K$, the feature aligns to the weights:

$$\mathbf{W}_{EX}\mathbf{h}_k^0 \propto \mathbf{W}_k.$$

In DEQ cases:

$$f(\mathbf{h}_k^0; \mathbf{W}_{DEQ}) \propto \mathbf{W}_k.$$

The theorem demonstrates that when the network training reaches its limit, i.e., when the loss function reaches its minimum, the \mathcal{NC} phenomenon emerges regardless of whether the chosen network is DEQ or explicit neural network. Besides, in certain scenarios, the lower bound of the loss function for DEQ is relatively smaller compared to explicit neural networks. More detailed proofs are in Appendix Section B.

4 Comparison under imbalanced setting

In this section, we analyze the performance differences between DEQ and explicit neural network on imbalanced datasets. We observe that, unlike in balanced scenarios, as long as certain conditions are met, the advantages of DEQ over explicit neural network become more pronounced on imbalanced datasets. And we provide theoretical evidence to support this phenomenon.

Suppose the total number of classes is K , with K_A being the number of majority classes and $K_B = K - K_A$ being the number of minority classes. Each majority class has n_A samples, and each minority class has n_B samples. The total number of samples is given by $N = K_A n_A + K_B n_B$. Note that $n_A > n_B$ with no requirement for K_A to be greater than K_B . We first start with the loss function, which can be partitioned into two components as follows:

$$\begin{aligned} \min_{\mathbf{W}, \tilde{\mathbf{W}}, \mathbf{H}^0} & \frac{K_A n_A}{N} \sum_{k=1}^{K_A} \sum_{i=1}^{n_A} \mathcal{L}(\mathbf{W} \tilde{\mathbf{W}} \mathbf{H}^0, \mathbf{y}_k) + \frac{K_B n_B}{N} \sum_{k=K_A+1}^{K_B} \sum_{i=1}^{n_B} \mathcal{L}(\mathbf{W} \tilde{\mathbf{W}} \mathbf{H}^0, \mathbf{y}_k), \\ \text{s.t. } & \tilde{\mathbf{W}} \in \{\mathcal{C}_{\text{EX}} \text{ or } \mathcal{C}_{\text{DEQ}}\}, \mathbf{w}_k, \mathbf{h}_{k,i} \in \mathcal{C}, \end{aligned} \quad (10)$$

where $\tilde{\mathbf{W}}$ represents the weights of Deep Equilibrium Models \mathbf{W}_{DEQ} and explicit neural network \mathbf{W}_{EX} . To analyze the \mathcal{NC} phenomenon, we present the results in the following theorem:

Theorem 4.1. (Neural Collapse under imbalanced settings on explicit neural networks and deep equilibrium models)

When the loss function reaches the minimum, then

$\mathcal{NC1}$: For $\forall k = 1, 2, \dots, K$ and $\forall i = 1, 2, \dots, n$:

$$\mathbf{W}_{\text{EX}} \mathbf{h}_{k,i}^0 = \mathbf{W}_{\text{EX}} \mathbf{h}_k^0,$$

where $\mathbf{h}_k^0 = \sum_{i \in \tau(i)} \mathbf{h}_{k,i}^0$. Similarly, if the model is DEQ, then

$$f(\mathbf{h}_{k,i}^0; \mathbf{W}_{\text{DEQ}}) = f(\mathbf{h}_k^0; \mathbf{W}_{\text{DEQ}}).$$

$\mathcal{NC2}$: Not exists, but the results of explicit neural network and DEQ can be compared:

Here we denote $(\mathbf{h}_k^0)^T \mathbf{h}_{k'}^0 = \mathbf{m}_{k,k'}$ and \mathbf{S} is a K -Simplex ETF, if

$$E_H < 2\mathbf{S}_{ij} - \mathbf{m}_{ij} < \frac{1}{1 - E_H}$$

is satisfied, the following inequality

$$\left\| (\mathbf{W}_{\text{EX}} \mathbf{H}^0)^T (\mathbf{W}_{\text{EX}} \mathbf{H}^0) - \mathbf{S} \right\|_F > \left\| f^T(\mathbf{H}^0; \mathbf{W}_{\text{DEQ}}) f(\mathbf{H}^0; \mathbf{W}_{\text{DEQ}}) - \mathbf{S} \right\|_F$$

holds.

$\mathcal{NC3}$: Similarly as $\mathcal{NC2}$, though it does not exist, the results can still be compared, when

$$\frac{E_H}{E_w + E_H} + E_H(1 - E_H) < 2$$

is satisfied, then the cosine distance satisfies:

$$\cos(f(\mathbf{h}_k; \mathbf{W}_{\text{DEQ}}), \mathbf{w}_k) / \cos(\mathbf{W}_{\text{EX}} \mathbf{h}_k, \mathbf{w}_k) > 1.$$

The detailed proof is in Appendix Section C.

Besides, the conclusion regarding the loss function is quite similar to that of Theorem B.3 under balanced settings. As analyzed in (43) and (44) in the Appendix, the lower bound of the loss function in DEQ is still lower than that in explicit neural network, where the performance of learned features is more evident in Figure 2, where we use t-SNE [53] and Gram matrix of features to describe the performance of two models. Although the phenomenon of $\mathcal{NC2}$ and $\mathcal{NC3}$ does not exist, we have discovered in Theorem 4.1 that under mild conditions, DEQ is superior in terms of \mathcal{NC} compared to

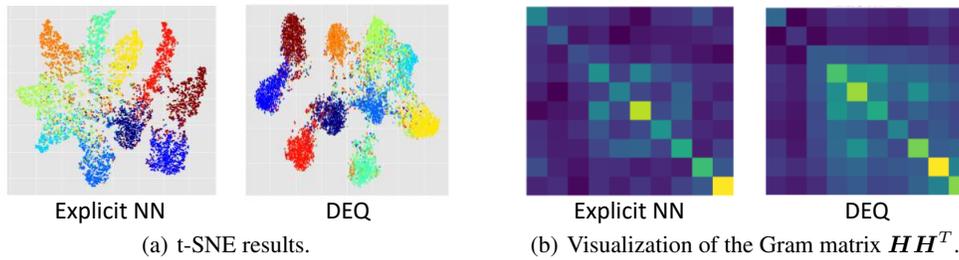


Figure 2: Under the imbalanced setting for CIFAR-10 with $K_A = 3$ and $R = 10$, the disparity in the learned features between Explicit Neural Networks (left) and DEQ (right).

explicit neural network. Notably, the conditions are easy to satisfy since E_H is generally very small in practice.

A crucial insight is that since DEQ undergoes multiple rounds of parameter adjustments for learning, it can be viewed as having an infinite number of layers, thus possessing greater representational capacity. As the network deepens, the iterative process of forward fixed-point may not necessarily reach the lowest threshold. Therefore, DEQ exhibits a certain degree of generalization for features in the minority class. Given the substantial feature differences among classes under an imbalanced dataset, the learned features by DEQ may demonstrate better adaptability to unseen categories. Consequently, compared to explicit neural network, DEQ tends to enhance performance.

Besides, due to the repeated iterations in solving the fixed-point iteration for some samples in the minority class with a small sample size, the model somewhat engages in multiple learning iterations for the features of samples in this class. This mitigates the impact of imbalanced samples to some extent. However, despite some improvements compared to the explicit neural network, DEQ still faces the issue of minority collapse. This conclusion is further validated in our subsequent experiments. Besides, to further discuss the situation of the dataset in terms of the degree of imbalance, we derived the following proposition:

Proposition 4.2. Denote $R = K_A n_A / N$. When the number of samples in the majority class becomes extremely large, i.e., $R \rightarrow 1$, the features of the two kinds of classes will become:

Majority classes:

$$\begin{aligned} \mathbf{W}_{EX} \mathbf{h}_{k,i}^0 &= \mathbf{W}_{EX} \mathbf{h}_k^0, \\ f(\mathbf{h}_{k,i}^0; \mathbf{W}_{DEQ}) &= f(\mathbf{h}_k^0; \mathbf{W}_{DEQ}), \end{aligned}$$

where $1 \leq k \leq K_A$ and $i \in \pi(k)$. Each feature collapses to K_A -Simplex ETF.

Minority classes:

$$\begin{aligned} \mathbf{w}_k &= \mathbf{0}, \\ \mathbf{W}_{EX} \mathbf{h}_{k,i}^0 &= f(\mathbf{h}_{k,i}^0; \mathbf{W}_{DEQ}) = \mathbf{0}, \end{aligned}$$

where $K_A + 1 \leq k \leq K$ and $i \in \pi(k)$.

Here, $\pi(k)$ refers to the samples that belong to the class k .

This situation is equivalent to having a balanced dataset in the majority class, while the minority class, due to its extremely small sample size, contributes almost nothing. In such an extreme scenario, the \mathcal{NC} performance of DEQ and the fully connected layer is nearly indistinguishable similar to Theorem 3.1. Both collapse on majority classes, resulting in a lack of learning features from minority classes meeting the results of (51) and (52). This aligns with the findings in [14], where they provide more specific bounds on the ratio K_A/K_B in their Theorem 5.

5 Experiments

In this section, we empirically conducted experiments to validate the correctness of the proposed theorems. Initially, we implemented DEQ on a balanced dataset and compared its \mathcal{NC} performance

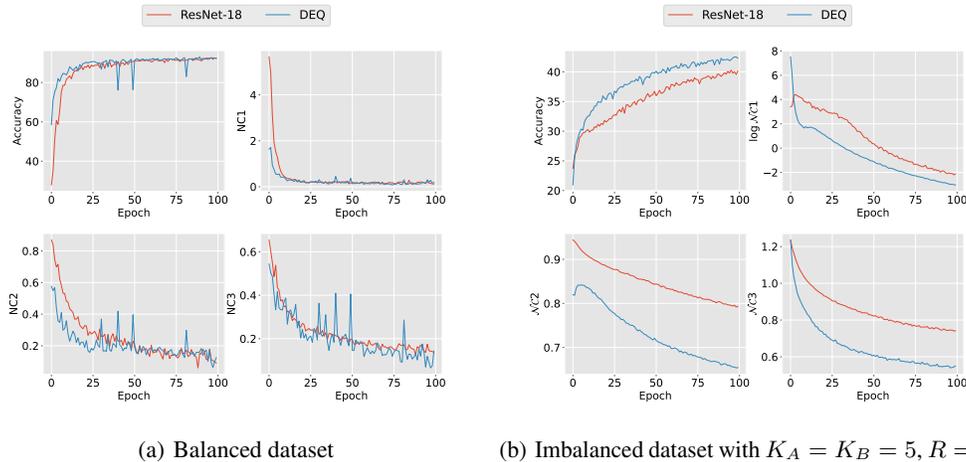


Figure 3: Comparison of accuracy and \mathcal{NC} phenomenon in training Cifar-10 dataset

with that of ResNet. Subsequently, for imbalanced datasets, we tested varying degrees of imbalance by manipulating the quantities of n_A and n_B , as well as K_A and K_B . The experimental results showed that, on imbalanced datasets, DEQ outperformed Explicit Neural Networks. This finding is consistent with the results reported in [4]. All experiments were implemented using PyTorch on NVIDIA Tesla A40 48GB.

5.1 Experiment setup

Without loss of generality, since any traditional neural network can be formulated as a DEQ, we use ResNet18 [22] as the backbone architecture here. As discussed earlier, to utilize the fixed point z^* learned by DEQ as the extracted feature, we formulate the last ResNet block into a DEQ format, while maintaining the remaining structure identical to ResNet. As mentioned in [5], training with DEQ can lead to instability issues. This is especially noticeable as training progresses, where some samples struggle to converge to a fixed point. To address this, in accordance with their setting, we implement the solver with a threshold ϵ set to 10^{-3} and introduce an early stopping mechanism. If convergence is not achieved within $T > 20$ iterations, we terminate the fixed-point iteration. Additionally, when facing problematic samples during fixed-point solving, we skip them to ensure training stability. During training, we set the learning rate to 1×10^{-4} and utilize stochastic gradient descent with a momentum of 0.9 and weight decay of 5×10^{-4} . Both E_W and E_H are set to 0.01. The training phase for each network consists of 100 epochs, with a batch size of 128. In this context, accuracy is assessed by averaging the results from the last 10 epochs and computing their standard deviation.

5.2 Performance under balanced conditions

By using the settings in (7) and (9), we compared the performance of DEQ and Explicit NN on Cifar-10 [30] and Cifar-100 [29] for validation, as shown in Figure 3(a). Their \mathcal{NC} performances remain comparable, i.e., DEQ achieves results similar to Explicit NN, corroborating the findings of Theorem 3.1. As for accuracy, from the

Table 1: Comparison of accuracy under balanced settings of Cifar-10 and Cifar-100

	Cifar-10	Cifar-100
Explicit NN	93.05 ± 0.17	64.35 ± 0.20
DEQ	93.23 ± 0.13	64.77 ± 0.36

results in the first column of Table 1, it can be observed that DEQ's accuracy is higher than that of the explicit layer, which aligns with Theorem B.3. However, the increase is only marginal due to the fact that the coefficients E_H and E_W act as scaling factors. Therefore, compared to explicit neural network, DEQ finds it challenging to achieve a significantly lower loss and, consequently, a substantial improvement. Moreover, Explicit NN performs well in fitting balanced datasets, so the accuracy of DEQ does not experience a significant boost in this context.

Here, we manually set the number of epochs to 100 to avoid potential instability issues with DEQ as training deepens. This is because DEQ can be challenging to reach the TPT (Terminal Phase of Training). As the number of parameters increases, achieving fixed-point convergence becomes more difficult, and even parameter explosion may occur. Under the current vanilla design, it is challenging to avoid such instability. Therefore, for a fair comparison, we apply the same training settings to both the implicit DEQ and the explicit neural network. The results in Figure 5 indicate that the test performance at 100 epochs is not significantly different from that at TPT. Since DEQ shares the same backbone as the corresponding explicit neural network, it can still demonstrate better \mathcal{NC} performance after reaching TPT in these cases.

Table 2: Test Accuracy on Cifar-10 and Cifar-100 Dataset with $K_A = 3$

	R	Cifar-10			Cifar-100		
		10	50	100	10	50	100
Explicit NN	overall	72.57±0.25	44.32±0.23	32.14±0.81	41.41±0.56	28.18±0.42	23.43±0.92
	majority	96.40±0.32	96.80±0.29	91.67±0.61	73.03±0.62	74.53±0.55	73.46±0.56
	minority	62.36±0.12	21.83±0.20	6.64±0.99	27.86±0.39	8.31±0.38	1.99±1.06
DEQ	overall	73.84±0.72	46.08±1.06	34.18±1.28	43.72±0.60	30.46±1.27	24.78±1.93
	majority	96.68±0.87	96.63±0.98	93.33±1.36	74.16±0.82	73.63±0.95	74.89±0.88
	minority	64.06±0.66	24.42±1.32	8.83±1.08	30.67±0.53	11.96±1.66	3.31±2.45

5.3 Performance under imbalanced conditions

We conducted experiments with varying configurations with different numbers of majority and minority classes and imbalance degrees. Assume the numbers of majority and minority classes are (K_A, K_B) with corresponding sample sizes (n_A, n_B) , the imbalance degree is denoted as $R = n_A/n_B$. We considered different setups for majority and minority class quantities, such as $(3, 7)$, $(5, 5)$, and $(7, 3)$. Additionally, we varied the ratio of sample quantities R between majority and minority classes with values of 10, 50 and 100. We also tested the phenomenon of \mathcal{NC} and accuracy on the Cifar-10 and Cifar-100 datasets, which own a total of 5000 images for each class. Specifically, when $R = 100$ and $(K_A, K_B) = (3, 7)$ for Cifar-10, the number of samples for all classes is $(5000, 5000, 5000, 50, 50, 50, 50, 50, 50, 50)$.

The results for $(K_A, K_B) = (3, 7)$ are shown in Table 2, where the test dataset owns the same distribution as the training dataset. We use “overall”, “majority”, and “minority” to represent the results across all categories, the majority class, and the minority class, respectively. We contrasted the difference in the training outcomes between the Explicit Neural Network and DEQ, and the superior performance of DEQ compared to Explicit Neural Network confirms DEQ’s higher learning potential. This suggests that DEQ can achieve a lower bound on its loss function. The experimental results indicate that DEQ consistently outperforms explicit neural network in accuracy during imbalanced training, aligning with Theorem 4.1. Specifically, we present the outcomes for $(K_A, K_B) = (5, 5)$ with $R = 100$ are depicted in Figure 3(b). The results strongly corroborate Theorem 4.1, affirming DEQ exhibits the same $\mathcal{NC}1$ phenomenon as an explicit neural network under these conditions. However, DEQ outperforms the explicit neural network in terms of $\mathcal{NC}2$ and $\mathcal{NC}3$. Additional experimental results with different parameters are detailed in Appendix Section D.

In addition to the stability considerations discussed in Section 5.2, we refrain from training for an extensive number of epochs due to the imbalance in the samples of the training set. This is because excessive learning rounds might cause the network parameters to predominantly capture information from the majority class, resulting in overfitting its features. This, in turn, diminishes the generalization of learning features from other classes, leading to marginal improvements in accuracy on the test set. As depicted in Figure 3(b), the model has already converged at this point. Moreover, limiting the number of training epochs helps to avoid the gradual instability in the learning process of DEQ.

6 Conclusion

In this study, we have systematically analyzed the representation of Deep Equilibrium Models (DEQ) and explicit neural networks under both balanced and imbalanced conditions using the phenomenon of Neural Collapse (\mathcal{NC}). Our theoretical analysis demonstrated that \mathcal{NC} is present in DEQ under

balanced conditions. Furthermore, in imbalanced settings, DEQ exhibited notable advantages over explicit neural networks, such as the convergence of extracted features to the vertices of a simplex equiangular tight frame and self-duality properties under mild conditions. These findings highlight the superior performance of DEQ in handling imbalanced datasets. Our experimental results in both balanced and imbalanced scenarios validate the theoretical insights. The current analysis is limited to simple imbalanced scenarios and the linear structure of DEQ models. Future work will expand on this foundation by exploring more general imbalanced scenarios and extending the analysis to more complex forms of DEQ models.

Acknowledgement

This work was supported by NSFC (No.62303319), Shanghai Sailing Program (22YF1428800), Shanghai Local College Capacity Building Program (23010503100), ShanghaiTech AI4S Initiative SHTAI4S202404, Shanghai Frontiers Science Center of Human-centered Artificial Intelligence (ShangHAI), MoE Key Laboratory of Intelligent Perception and Human-Machine Collaboration (ShanghaiTech University) and Shanghai Engineering Research Center of Intelligent Vision and Imaging.

References

- [1] A. Agrawal, B. Amos, S. Barratt, S. Boyd, S. Diamond, and Z. Kolter. Differentiable convex optimization layers. In *Advances in Neural Information Processing Systems*, 2019.
- [2] Brandon Amos and J Zico Kolter. Optnet: Differentiable optimization as a layer in neural networks. In *International Conference on Machine Learning*, pages 136–145. PMLR, 2017.
- [3] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. Deep equilibrium models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [4] Shaojie Bai, Vladlen Koltun, and J. Zico Kolter. Multiscale deep equilibrium models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [5] Shaojie Bai, Vladlen Koltun, and J. Zico Kolter. Stabilizing equilibrium models by jacobian regularization. In *International Conference on Machine Learning (ICML)*, 2021.
- [6] Mathieu Blondel, Quentin Berthet, Marco Cuturi, Roy Frostig, Stephan Hoyer, Felipe Llinares-López, Fabian Pedregosa, and Jean-Philippe Vert. Efficient and modular implicit differentiation. *Advances in neural information processing systems*, 35:5230–5242, 2022.
- [7] Charles G Broyden. A class of methods for solving nonlinear simultaneous equations. *Mathematics of computation*, 19(92):577–593, 1965.
- [8] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. *Advances in neural information processing systems*, 32, 2019.
- [9] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- [10] Hien Dang, Tan Nguyen, Tho Tran, Hung Tran, and Nhat Ho. Neural collapse in deep linear network: From balanced to imbalanced data. *arXiv preprint arXiv:2301.00437*, 2023.
- [11] Shutong Ding, Tianyu Cui, Jingya Wang, and Ye Shi. Two sides of the same coin: Bridging deep equilibrium models and neural ODEs via homotopy continuation. *Advances in Neural Information Processing Systems*, 36, 2024.
- [12] Justin Domke. Generic methods for optimization-based modeling. In *Artificial Intelligence and Statistics*, pages 318–326. PMLR, 2012.
- [13] Laurent El Ghaoui, Fangda Gu, Bertrand Travacca, Armin Askari, and Alicia Tsai. Implicit deep learning. *SIAM Journal on Mathematics of Data Science*, 3(3):930–958, 2021.
- [14] Cong Fang, Hangfeng He, Qi Long, and Weijie J Su. Exploring deep neural networks via layer-peeled model: Minority collapse in imbalanced training. *Proceedings of the National Academy of Sciences*, 118(43):e2103091118, 2021.

- [15] Samy Wu Fung, Howard Heaton, Qiuwei Li, Daniel McKenzie, Stanley Osher, and Wotao Yin. Jfb: Jacobian-free backpropagation for implicit networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 6648–6656, 2022.
- [16] Tomer Galanti, András György, and Marcus Hutter. On the role of neural collapse in transfer learning. *arXiv preprint arXiv:2112.15121*, 2021.
- [17] Tianxiang Gao, Xiaokai Huo, Hailiang Liu, and Hongyang Gao. Wide neural networks as gaussian processes: Lessons from deep equilibrium models. *arXiv preprint arXiv:2310.10767*, 2023.
- [18] Zhengyang Geng, Xin-Yu Zhang, Shaojie Bai, Yisen Wang, and Zhouchen Lin. On training implicit models. *Advances in Neural Information Processing Systems*, 34:24247–24260, 2021.
- [19] Davis Gilton, Gregory Ongie, and Rebecca Willett. Deep equilibrium architectures for inverse problems in imaging. *IEEE Transactions on Computational Imaging*, 7:1123–1133, 2021.
- [20] Stephen Gould, Richard Hartley, and Dylan Campbell. Deep declarative networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(8):3988–4004, 2021.
- [21] Fangda Gu, Heng Chang, Wenwu Zhu, Somayeh Sojoudi, and Laurent El Ghaoui. Implicit graph neural networks. *Advances in Neural Information Processing Systems*, 33:11984–11995, 2020.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [23] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [24] Yuhao Huang, Qingsong Wang, Akwum Onwunta, and Bao Wang. Efficient score matching with deep equilibrium layers. In *The Twelfth International Conference on Learning Representations*, 2024.
- [25] Like Hui, Mikhail Belkin, and Preetum Nakkiran. Limitations of neural collapse for understanding generalization in deep learning. *arXiv preprint arXiv:2202.08384*, 2022.
- [26] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.
- [27] Jiachen Jiang, Jinxin Zhou, Peng Wang, Qing Qu, Dustin Mixon, Chong You, and Zhihui Zhu. Generalized neural collapse for a large number of classes. *arXiv preprint arXiv:2310.05351*, 2023.
- [28] Vignesh Kothapalli. Neural collapse: A review on modelling principles and generalization. *Transactions on Machine Learning Research*, 2023.
- [29] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical Report 0, University of Toronto, Toronto, Ontario, 2009.
- [30] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced research). URL <http://www.cs.toronto.edu/kriz/cifar.html>, 5(4):1, 2010.
- [31] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.
- [32] Jaehoon Lee, Yasaman Bahri, Roman Novak, Samuel S Schoenholz, Jeffrey Pennington, and Jascha Sohl-Dickstein. Deep neural networks as gaussian processes. *arXiv preprint arXiv:1711.00165*, 2017.
- [33] Mingchen Li, Xuechen Zhang, Christos Thrampoulidis, Jiasi Chen, and Samet Oymak. Autobalance: Optimized loss functions for imbalanced data. *Advances in Neural Information Processing Systems*, 34:3163–3177, 2021.
- [34] Mingjie Li, Yisen Wang, and Zhouchen Lin. Cerdeq: Certifiable deep equilibrium model. In *International Conference on Machine Learning*, pages 12998–13013. PMLR, 2022.

- [35] Zenan Ling, Longbo Li, Zhanbo Feng, YIXUAN ZHANG, Feng Zhou, Robert C Qiu, and Zhenyu Liao. Deep equilibrium models are almost equivalent to not-so-deep explicit models for high-dimensional gaussian mixtures. In *Forty-first International Conference on Machine Learning*, 2024.
- [36] Zenan Ling, Xingyu Xie, Qiuhaio Wang, Zongpeng Zhang, and Zhouchen Lin. Global convergence of over-parameterized deep equilibrium models. In *International Conference on Artificial Intelligence and Statistics*, pages 767–787. PMLR, 2023.
- [37] Xuantong Liu, Jianfeng Zhang, Tianyang Hu, He Cao, Yuan Yao, and Lujia Pan. Inducing neural collapse in deep long-tailed learning. In *International Conference on Artificial Intelligence and Statistics*, pages 11534–11544. PMLR, 2023.
- [38] Jonathan Lorraine, Paul Vicol, and David Duvenaud. Optimizing millions of hyperparameters by implicit differentiation. In *International conference on artificial intelligence and statistics*, pages 1540–1552. PMLR, 2020.
- [39] Song Mei, Andrea Montanari, and Phan-Minh Nguyen. A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671, 2018.
- [40] Dustin G Mixon, Hans Parshall, and Jianzong Pi. Neural collapse with unconstrained features. *arXiv preprint arXiv:2011.11619*, 2020.
- [41] Vishal Monga, Yuelong Li, and Yonina C Eldar. Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing. *IEEE Signal Processing Magazine*, 38(2):18–44, 2021.
- [42] Vardan Papyan, XY Han, and David L Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40):24652–24663, 2020.
- [43] Ashwini Pople, Zhengyang Geng, and J Zico Kolter. Deep equilibrium approaches to diffusion models. *Advances in Neural Information Processing Systems*, 35:37975–37990, 2022.
- [44] Zaccharie Ramzi, Pierre Ablin, Gabriel Peyré, and Thomas Moreau. Test like you train in implicit deep learning. *arXiv preprint arXiv:2305.15042*, 2023.
- [45] Akshay Rangamani, Marius Lindegaard, Tomer Galanti, and Tomaso A Poggio. Feature learning in deep classifiers through intermediate neural collapse. In *International Conference on Machine Learning*, pages 28729–28745. PMLR, 2023.
- [46] Jiawei Ren, Cunjun Yu, Xiao Ma, Haiyu Zhao, Shuai Yi, et al. Balanced meta-softmax for long-tailed visual recognition. *Advances in neural information processing systems*, 33:4175–4186, 2020.
- [47] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *International conference on machine learning*, pages 4334–4343. PMLR, 2018.
- [48] Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Xiaojiang Chen, and Xin Wang. A comprehensive survey of neural architecture search: Challenges and solutions. *ACM Computing Surveys (CSUR)*, 54(4):1–34, 2021.
- [49] Haixiang Sun, Ye Shi, Jingya Wang, Hoang Duong Tuan, H Vincent Poor, and Dacheng Tao. Alternating differentiation for optimization layers. *arXiv preprint arXiv:2210.01802*, 2022.
- [50] Christos Thrampoulidis, Ganesh Ramachandra Kini, Vala Vakilian, and Tina Behnia. Imbalance trouble: Revisiting neural-collapse geometry. *Advances in Neural Information Processing Systems*, 35:27225–27238, 2022.
- [51] Tom Tirer and Joan Bruna. Extended unconstrained features model for exploring deep neural collapse. In *International Conference on Machine Learning*, pages 21478–21505. PMLR, 2022.
- [52] Russell Tsuchida, Suk Yee Yong, Mohammad Ali Armin, Lars Petersson, and Cheng Soon Ong. Declarative nets that are equilibrium models. In *International Conference on Learning Representations*, 2022.
- [53] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(86):2579–2605, 2008.

- [54] Yuchao Wang, Jingjing Fei, Haochen Wang, Wei Li, Tianpeng Bao, Liwei Wu, Rui Zhao, and Yujun Shen. Balancing logit variation for long-tailed semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19561–19573, 2023.
- [55] Colin Wei and J Zico Kolter. Certified robustness for deep equilibrium models via interval bound propagation. In *International Conference on Learning Representations*, 2022.
- [56] Ezra Winston and J Zico Kolter. Monotone operator equilibrium networks. *Advances in neural information processing systems*, 33:10718–10728, 2020.
- [57] Liang Xie, Yibo Yang, Deng Cai, and Xiaofei He. Neural collapse inspired attraction–repulsion-balanced loss for imbalanced learning. *Neurocomputing*, 527:60–70, 2023.
- [58] Xingyu Xie, Qiuhaio Wang, Zenan Ling, Xia Li, Guangcan Liu, and Zhouchen Lin. Optimization induced equilibrium networks: An explicit optimization perspective for understanding equilibrium models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3):3604–3616, 2022.
- [59] Yibo Yang, Shixiang Chen, Xiangtai Li, Liang Xie, Zhouchen Lin, and Dacheng Tao. Inducing neural collapse in imbalanced learning: Do we really need a learnable classifier at the end of deep neural network? *Advances in Neural Information Processing Systems*, 35:37991–38002, 2022.
- [60] Zonghan Yang, Tianyu Pang, and Yang Liu. A closer look at the adversarial robustness of deep equilibrium models. *Advances in Neural Information Processing Systems*, 35:10448–10461, 2022.
- [61] Can Yaras, Peng Wang, Zhihui Zhu, Laura Balzano, and Qing Qu. Neural collapse with normalized features: A geometric analysis over the riemannian manifold. *Advances in neural information processing systems*, 35:11547–11560, 2022.
- [62] Yifan Zhang, Bingyi Kang, Bryan Hooi, Shuicheng Yan, and Jiashi Feng. Deep long-tailed learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
- [63] Zhisheng Zhong, Jiequan Cui, Yibo Yang, Xiaoyang Wu, Xiaojuan Qi, Xiangyu Zhang, and Jiaya Jia. Understanding imbalanced semantic segmentation through neural collapse. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19550–19560, 2023.
- [64] Jinxin Zhou, Xiao Li, Tianyu Ding, Chong You, Qing Qu, and Zhihui Zhu. On the optimization landscape of neural collapse under mse loss: Global optimality with unconstrained features. In *International Conference on Machine Learning*, pages 27179–27202. PMLR, 2022.
- [65] Jianggang Zhu, Zheng Wang, Jingjing Chen, Yi-Ping Phoebe Chen, and Yu-Gang Jiang. Balanced contrastive learning for long-tailed visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6908–6917, 2022.
- [66] Zhihui Zhu, Tianyu DING, Jinxin Zhou, Xiao Li, Chong You, Jeremias Sulam, and Qing Qu. A geometric analysis of neural collapse with unconstrained features. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021.

Appendix Contents

A	Evaluation metrics of \mathcal{NC}	14
B	Proof under balanced setting	15
B.1	Problem definition	15
B.2	\mathcal{NC} analysis	17
B.2.1	\mathcal{NC} proof in Explicit neural networks	17
B.2.2	\mathcal{NC} proof in DEQ	18
C	Proof under imbalanced learning	20
C.1	Lower bound of the loss function	20
C.2	\mathcal{NC} Analysis	23
D	More experiments	26

A Evaluation metrics of \mathcal{NC}

Followed by the settings of [51] and [10], the measurement of \mathcal{NC} are set as follow:

Let $\mathbf{h}_k \triangleq \frac{1}{n_k} \sum_{i=1}^{n_k} \mathbf{h}_{k,i}$ represent the average of all features within class k and these K classes collectively constitute the average matrix $\bar{\mathbf{H}} = [\mathbf{h}_1, \dots, \mathbf{h}_K]$. Besides, The global average is defined as $\mathbf{h}_G \triangleq \frac{1}{K} \sum_{i=1}^K \mathbf{h}_k$. Subsequently, the within-class and between-class covariances can be calculated as:

$$\begin{aligned} \Sigma_W &\triangleq \frac{1}{N} \sum_{k=1}^K \sum_{i=1}^{n_k} (\mathbf{h}_{k,i} - \mathbf{h}_k)(\mathbf{h}_{k,i} - \mathbf{h}_k)^T, \\ \Sigma_B &\triangleq \frac{1}{K} \sum_{k=1}^K (\mathbf{h}_k - \bar{\mathbf{h}}_G)(\mathbf{h}_k - \bar{\mathbf{h}}_G)^T. \end{aligned} \quad (11)$$

$\mathcal{NC}1$ measures the variation of features with-in the same class:

$$\mathcal{NC}1 = \frac{1}{K} \text{tr} \left(\Sigma_W \Sigma_B^\dagger \right), \quad (12)$$

where Σ_B^\dagger denotes the pseudo-inverse of Σ_B .

$\mathcal{NC}2$ measures similarity between the mean of learned last-layer features $\bar{\mathbf{H}}$ and the structure of Simplex ETF:

$$\mathcal{NC}2 = \left\| \frac{\bar{\mathbf{H}}^T \bar{\mathbf{H}}}{\|\bar{\mathbf{H}}^T \bar{\mathbf{H}}\|_F} - \frac{1}{K-1} \left(\mathbf{I}_K - \frac{1}{K} \mathbf{1}_K \mathbf{1}_K^T \right) \right\|_F. \quad (13)$$

$\mathcal{NC}3$ measures similarity of the last-layer feature $\bar{\mathbf{H}}$ and weights of classifier \mathbf{W} :

$$\mathcal{NC}3 = \left\| \frac{\mathbf{W}}{\|\mathbf{W}\|_F} - \frac{\bar{\mathbf{H}}}{\|\bar{\mathbf{H}}\|_F} \right\|. \quad (14)$$

Additionally, it is worth noting that all above \mathcal{NC} criteria are exclusively based on the training set. This is because our focus is solely on analyzing learning performance on imbalanced datasets, and generalization is not a primary concern.

B Proof under balanced setting

B.1 Problem definition

As different layers in the neural network introduce complexity, the optimization problem is non-convex, and KKT conditions do not guarantee global optimality. Therefore, we consider applying inequality relaxation to the joint optimization problem, obtaining a lower bound for the loss function. By determining the conditions under which the equality holds, we can derive the requirements for the \mathcal{NC} phenomenon. This analysis assumes a balanced setting, where all $\#\tau(k) = n_1 = n_2 = \dots = n_K = N/K$.

We considered the fully connected layers (explicit) and Deep Equilibrium Models (implicit) under the balanced settings respectively, and then derived the detailed proof.

(Fully Connected Layers)

$$\begin{aligned} \min_{\mathbf{W}, \mathbf{W}_{\text{EX}}, \mathbf{H}} \quad & \frac{1}{N} \sum_{k=1}^K \sum_{i=1}^{n_k} \mathcal{L}(\mathbf{W} \mathbf{W}_{\text{EX}} \mathbf{h}_{k,i}^0, \mathbf{y}_k) \\ \text{s.t.} \quad & \mathbf{h}_{k,i} = \mathbf{W}_{\text{EX}} \mathbf{h}_{k,i}^0, \\ & \|\mathbf{W}_{\text{EX}}\|_F \leq E_H, \\ & \frac{1}{K} \sum_{k=1}^K \|\mathbf{w}_k\|^2 \leq E_W, \\ & \frac{1}{K} \sum_{k=1}^K \frac{1}{n_k} \sum_{i=1}^{n_k} \|\mathbf{h}_{k,i}\|^2 \leq E_H, \end{aligned}$$

(Deep Equilibrium Models)

$$\begin{aligned} \min_{\mathbf{W}, \mathbf{W}_{\text{DEQ}}, \mathbf{z}_{k,i}^*} \quad & \frac{1}{N} \sum_{k=1}^K \sum_{i=1}^{n_k} \mathcal{L}(\mathbf{W} \mathbf{z}_{k,i}^*, \mathbf{y}_k) \\ \text{s.t.} \quad & \mathbf{z}_{k,i}^* = f(\mathbf{h}_{k,i}^0; \mathbf{W}_{\text{DEQ}}), \\ & \|\mathbf{W}_{\text{DEQ}}\|_F \leq E_H, \\ & \frac{1}{K} \sum_{k=1}^K \|\mathbf{w}_k\|^2 \leq E_W, \\ & \frac{1}{K} \sum_{k=1}^K \frac{1}{n_k} \sum_{i=1}^{n_k} \|\mathbf{z}_{k,i}^*\|^2 \leq E_H. \end{aligned}$$

Note that here $n_1 = n_2 = \dots = n_k = n$, and f represents the form of Linear DEQ, where we will use $f(\mathbf{x}; \mathbf{W}_{\text{DEQ}}) = \sum_{i=1}^{\infty} \mathbf{W}_{\text{DEQ}}^i \mathbf{x}$ for representation in the following proofs.

In a classification task, cross-entropy loss $\mathcal{L}(\mathbf{W} \mathbf{h}_{k,i}, \mathbf{y}_k)$ is regarded as the final loss function. Drawing inspiration from [14], our initial efforts revolve around organizing and simplifying the log function to distinguish the logit in class k from other classes.

First consider the following lemma:

Lemma B.1. *Let there be K variables $\delta_1, \delta_2, \dots, \delta_K$, and the logit of each variable δ_k satisfies the inequality:*

$$\log \left(\frac{\delta_k}{\sum_{k=1}^K \delta_k} \right) \leq M_1 \left(\log \delta_k - \frac{1}{K-1} \sum_{k' \neq k}^K \log \delta_{k'} \right) + M_2, \quad (15)$$

where M_1 and M_2 are predefined constants.

Proof. Split the sum in the denominator and sequentially introduce weights for each term. Here, define K coefficients such that their sum is 1. Therefore, we have:

$$\frac{C_1}{C_1 + C_2} + \underbrace{C_3 + \dots + C_3}_{K-1} = 1, \quad (16)$$

that is $C_3 = \frac{C_2}{(K-1)(C_1+C_2)}$. Therefore, by Jensen's inequality, we can derive:

$$\begin{aligned} \log \left(\delta_k / \sum_{k'=1}^K \delta_{k'} \right) &= \log \delta_k - \log \left(\sum_{k'=1}^K \delta_{k'} \right) \\ &= \log \delta_k - \log \left(\frac{C_1}{C_1+C_2} \frac{(C_1+C_2)\delta_k}{C_1} + C_3 \sum_{k' \neq k}^K \frac{\delta_{k'}}{C_3} \right) \\ &\leq \log \delta_k - \frac{C_1}{C_1+C_2} \log \left(\frac{(C_1+C_2)\delta_k}{C_1} \right) - C_3 \sum_{k' \neq k}^K \log \frac{\delta_{k'}}{C_3} \\ &= M_1 \left(\log \delta_k - \frac{1}{K-1} \sum_{k' \neq k}^K \log \delta_{k'} \right) + M_2, \end{aligned} \tag{17}$$

where $M_1 = \frac{C_2}{C_1+C_2}$ and $M_2 = \frac{C_2}{C_1+C_2} \log C_3 - \frac{C_1}{C_1+C_2} \log \left(\frac{C_1+C_2}{C_1} \right)$. Therefore the lemma is proved. \square

Remark B.2. When $C_2/C_1 = \frac{1}{K-1} \exp \left(\log \delta_k - \frac{1}{K-1} \sum_{k' \neq k}^K \log \delta_{k'} \right)$, the right term of the inequality in lemma B.1 reaches its maximum.

Proof. Let $\mathcal{M} = M_1 \left(\log \delta_k - \frac{1}{K-1} \sum_{k' \neq k}^K \log \delta_{k'} \right) + M_2$ in Lemma B.1. Then, upon computing the derivatives of C_1 and C_2 , we obtain:

$$\begin{aligned} \frac{\partial \mathcal{M}}{\partial C_1} &= \frac{1}{(C_1+C_2)^2} \left(-C_2 \frac{\mathcal{M}-M_2}{M_1} + C_2 \log \frac{(K-1)C_2}{C_1} \right) \\ \frac{\partial \mathcal{M}}{\partial C_2} &= \frac{1}{(C_1+C_2)^2} \left(C_1 \frac{\mathcal{M}-M_2}{M_1} - C_1 \log \frac{(K-1)C_2}{C_1} \right). \end{aligned} \tag{18}$$

Combining these two equations yields the conclusion. \square

Next, we substitute the result of each logit into the lemma B.1, from which we can derive:

$$\begin{aligned} \mathcal{L} &= -\frac{1}{N} \sum_{k=1}^K \sum_{i=1}^n \mathbf{y}_{k,i} \log \frac{\exp(\mathbf{w}_{k,i}^T \mathbf{h}_i)}{\sum_{k'=1}^K \exp(\mathbf{w}_{k',i}^T \mathbf{h}_i)} \\ &\geq \frac{C_1}{(C_1+C_2)N(K-1)} \sum_{i=1}^n \left[\left(\sum_{k=1}^K \mathbf{h}_{k,i} \right)^T \left(\sum_{k=1}^K \mathbf{w}_k \right) - K \sum_{k=1}^K \mathbf{h}_{k,i}^T \mathbf{w}_k \right] + C_4 \\ &= \frac{C_1 K}{(C_1+C_2)N(K-1)} \sum_{k=1}^K \sum_{i=1}^n (\bar{\mathbf{h}}_i - \mathbf{h}_{k,i})^T \mathbf{w}_k + C_4 \\ &\geq \frac{C_1}{(C_1+C_2)N(K-1)} \left(-\frac{K}{2} \sum_{k=1}^K \sum_{i=1}^n \|\bar{\mathbf{h}}_i - \mathbf{h}_{k,i}\|^2 / C_5 - \frac{C_5 N}{2} \sum_{k=1}^K \|\mathbf{w}_k\|^2 \right) + C_4, \end{aligned} \tag{19}$$

where the second inequality applies Mean Inequalities, and $C_4 = \frac{C_2}{N(C_1+C_2)} \log C_3 - \frac{C_1}{N(C_1+C_2)} \log \left(\frac{C_1+C_2}{C_1} \right)$. For convenience, we denote $\tilde{\mathcal{L}} = -\frac{K}{2} \sum_{k=1}^K \sum_{i=1}^n \|\bar{\mathbf{h}}_i - \mathbf{h}_{k,i}\|^2 / C_5 - \frac{C_5 N}{2} \sum_{k=1}^K \|\mathbf{w}_k\|^2$

$\frac{C_5 N}{2} \sum_{k=1}^K \|\mathbf{w}_k\|^2$. As the corresponding constraints have already been added in (6), specifically the constraint $\sum_{k=1}^K \|\mathbf{w}_k\|^2 \leq E_W$, our focus shifts to discussing the situation concerning the first term. Since it represents the features of the final layer, we separately explore the differences in its extraction when using DEQ and fully connected layers. First suppose the extracted feature by the backbone is \mathbf{h}^0 .

B.2 \mathcal{NC} analysis

We separately discuss the representation of \mathcal{NC} in the cases of Explicit NN and DEQ, and compare the lower bounds of the loss function.

B.2.1 \mathcal{NC} proof in Explicit neural networks

For convenience, we assume there is only one layer in the feature extractor, that is, $\mathbf{h} = \mathbf{W}_{\text{EX}} \mathbf{h}^0$, then the first term in $\tilde{\mathcal{L}}$ becomes:

$$\begin{aligned} -\frac{K}{2} \sum_{k=1}^K \sum_{i=1}^n \|\bar{\mathbf{h}}_i - \mathbf{h}_{k,i}\|^2 &= -\frac{K}{2} \sum_{k=1}^K \sum_{i=1}^n \|\mathbf{W}_{\text{EX}}(\bar{\mathbf{h}}_i^0 - \mathbf{h}_{k,i}^0)\|^2 \\ &\geq -\frac{K}{4} \sum_{k=1}^K \sum_{i=1}^n \left(\|\mathbf{W}_{\text{EX}}\|_F^2 + \|\bar{\mathbf{h}}_i^0 - \mathbf{h}_{k,i}^0\|^2 \right). \end{aligned} \quad (20)$$

Substituting them into the loss function (19), we can observe that:

$$\begin{aligned} \tilde{\mathcal{L}} &\geq -\frac{NK}{4C_5} \|\mathbf{W}_{\text{EX}}\|_F^2 - \frac{K}{4C_5} \sum_{k=1}^K \sum_{i=1}^n \|\bar{\mathbf{h}}_i^0 - \mathbf{h}_{k,i}^0\|^2 - \frac{C_5 NK}{2} E_W \\ &= -\frac{K^2}{4C_5} \sum_{i=1}^n \frac{1}{K} \sum_{k=1}^K \left(\|\mathbf{h}_{k,i}^0\|^2 - \|\bar{\mathbf{h}}_i^0\|^2 \right) - \frac{NK}{4C_5} \|\mathbf{W}_{\text{EX}}\|_F^2 - \frac{C_5 NK}{2} E_W \\ &\geq -\frac{KN}{4C_5} E_H - \frac{C_5 KN}{2} E_W + \frac{K^2}{4C_5} \sum_{i=1}^n \|\bar{\mathbf{h}}_i^0\|^2 - \frac{NK}{4C_5} E_H. \end{aligned} \quad (21)$$

To acquire the lower bound of the loss function, we assign the value $C_5 = \sqrt{E_H/E_W}$, the lower bound becomes:

$$\inf \mathcal{L}_{\text{EX}} = -\frac{C_1 K}{(C_1 + C_2)(K - 1)} \sqrt{E_W E_H} + C_4. \quad (22)$$

Furthermore, the condition $\|\bar{\mathbf{h}}_i^0\|^2 = 0$ should also be satisfied, indicating that the average of the features for the i -th sample, $\frac{1}{K} \sum_{k=1}^K \mathbf{h}_{k,i}^0$, is equal to zero.

The satisfaction conditions for the inequalities include the following:

- In Eq. (19): The first inequality becomes equality when

$$\frac{(C_1 + C_2) \mathbf{h}_{k,i}^T \mathbf{w}_k}{C_1} = \frac{\mathbf{h}_{k,i}^T \mathbf{w}_{k'}}{C_3}, \quad (23)$$

that is,

$$\mathbf{h}_{k,i}^T \mathbf{w}_k = \mathbf{h}_{k,i}^T \mathbf{w}_{k'} + \log \left(\frac{C_1(K - 1)}{C_2} \right). \quad (24)$$

The second inequality is reduced to equality when $\bar{\mathbf{h}}_i - \mathbf{h}_{k,i} = -C_5 \mathbf{w}_k$.

- In Eq. (20): $\|\mathbf{W}_{\text{EX}}\|_F^2 = \sum_{k=1}^K \sum_{i=1}^n \|\bar{\mathbf{h}}_i^0 - \mathbf{h}_{k,i}^0\|^2$.

- In Eq. (21): When the following condition $\frac{1}{K} \sum_{k=1}^K \|\mathbf{w}_k\|^2 = E_W$ and $\|\mathbf{W}_{\text{EX}}\|_F^2 = \frac{1}{K} \sum_{k=1}^K \sum_{i=1}^{n_k} \|\mathbf{h}_{k,i}^0\|^2 = E_H$ holds, the inequality was reduced to equality.

Since $\|\bar{\mathbf{h}}_i\|^2 = 0$, it follows that $\|\bar{\mathbf{h}}_i\|^2 = \|\mathbf{W}_{\text{EX}} \bar{\mathbf{h}}_i^0\|^2 = 0$. Combined with the condition $\frac{1}{K} \sum_{k=1}^K \|\mathbf{w}_k\|^2 = E_W$ and $\frac{1}{K} \sum_{k=1}^K \sum_{i=1}^{n_k} \|\mathbf{h}_{k,i}\|^2 = E_H$, therefore, $\mathbf{h}_k = \mathbf{h}_{k,i}$, for $\forall k$, that is, $\mathcal{NC}1$ is proved.

Consequently, $\mathbf{h}_{k,i} = C_5 \mathbf{w}_k$, demonstrating the validity of $\mathcal{NC}3$.

For $\mathcal{NC}2$, since

$$\begin{aligned} \sqrt{E_H/E_W} \|\mathbf{w}_k\|^2 &= \mathbf{h}_k \mathbf{w}_k = \mathbf{h}_k \mathbf{w}_{k'} + \log\left(\frac{C_1(K-1)}{C_2}\right) = \mathbf{W}_{\text{EX}} \mathbf{h}_k^0 \mathbf{W}_{k'} + \log\left(\frac{C_1(K-1)}{C_2}\right), \\ \sqrt{E_H/E_W} \|\mathbf{w}_{k'}\|^2 &= \mathbf{h}_{k'} \mathbf{w}_{k'} = \mathbf{h}_{k'} \mathbf{w}_k + \log\left(\frac{C_1(K-1)}{C_2}\right) = \mathbf{W}_{\text{EX}} \mathbf{h}_{k'}^0 \mathbf{W}_k + \log\left(\frac{C_1(K-1)}{C_2}\right) \end{aligned} \quad (25)$$

holds, by the equality conditions, $\|\mathbf{w}_k\|^2 = \|\mathbf{w}_{k'}\|^2 = E_W$.

Further, $\sum_{k=1}^K \mathbf{h}_k \mathbf{w}_{k'} = \sum_{k=1}^K \mathbf{W}_{\text{EX}} \mathbf{h}_k^0 \mathbf{w}_{k'} = 0$, as $\mathbf{h}_k \mathbf{w}_k = \sqrt{E_W E_H}$, so $\mathbf{h}_k \mathbf{w}_{k'} = -\frac{\sqrt{E_W E_H}}{N-1}$.

Therefore, the $\mathcal{NC}2$ condition satisfies:

$$\mathbf{W} \mathbf{W}^T = \sqrt{E_W/E_H} \mathbf{W} \mathbf{H} = \frac{K E_W}{K-1} \left(\mathbf{1}_K - \frac{1}{k} \mathbf{1}_K \mathbf{1}_K^T \right). \quad (26)$$

B.2.2 \mathcal{NC} proof in DEQ

In the blocks for feature extraction, DEQ can be referred as a mapping from the features by backbone to the output $\mathbf{h}^0 \rightarrow \mathbf{h}^*$, which can be directly solved using the implicit equation:

$$\mathbf{h}^* = f(\mathbf{W}_{\text{DEQ}}; \mathbf{h}^0) = \sum_{i=1}^{\infty} \mathbf{W}_{\text{DEQ}}^i \mathbf{h}^0. \quad (27)$$

Similar as the explicit case, start with the term:

$$-\frac{K}{2} \sum_{k=1}^K \sum_{i=1}^n \|\bar{\mathbf{h}}_i - \mathbf{h}_{k,i}\|^2 = \frac{K}{2} \sum_{k=1}^K \sum_{i=1}^n \left\| \sum_{j=0}^{\infty} \mathbf{W}_{\text{DEQ}}^j (\bar{\mathbf{h}}_i^0 - \mathbf{h}_{k,i}^0) \right\|^2. \quad (28)$$

Since the Neumann series can be regarded as a recursive procedure, denote $\mathcal{G}_{k,i}^j = \sum_{j'=0}^j \mathbf{W}_{\text{DEQ}}^{j'} (\bar{\mathbf{h}}_i^0 - \mathbf{h}_{k,i}^0)$ ($j = 0, 1, \dots, \infty$), therefore $\mathcal{G}_{k,i}^j = \mathbf{W}_{\text{DEQ}} \mathcal{G}_{k,i}^{j-1} + (\bar{\mathbf{h}}_i^0 - \mathbf{h}_{k,i}^0)$.

$$\begin{aligned} -\frac{K}{2} \sum_{k=1}^K \sum_{i=1}^n \|\mathcal{G}_{k,i}^j\|^2 &= \frac{K}{2} \sum_{k=1}^K \sum_{i=1}^n \left\| \mathbf{W}_{\text{DEQ}} \mathcal{G}_{k,i}^{j-1} + (\bar{\mathbf{h}}_i^0 - \mathbf{h}_{k,i}^0) \right\|^2 \\ &\geq -\frac{K}{2} \sum_{k=1}^K \sum_{i=1}^n \|\mathbf{W}_{\text{DEQ}} \mathcal{G}_{k,i}^{j-1}\|^2 - \frac{K}{2} \sum_{k=1}^K \sum_{i=1}^n \|\bar{\mathbf{h}}_i^0 - \mathbf{h}_{k,i}^0\|^2 \\ &\geq -\frac{K}{4} \sum_{k=1}^K \sum_{i=1}^n \left(\|\mathbf{W}_{\text{DEQ}}\|_F^2 + \|\mathcal{G}_{k,i}^{j-1}\|^2 \right) - \frac{K}{2} \sum_{k=1}^K \sum_{i=1}^n \|\bar{\mathbf{h}}_i^0 - \mathbf{h}_{k,i}^0\|_F^2. \end{aligned} \quad (29)$$

Continuing the recursion, we can obtain:

$$-\frac{1}{2} \|\mathcal{G}_{k,i}^j\|^2 \geq -\left(\frac{1}{2}\right)^{j+1} \|\mathcal{G}_{k,i}^0\|^2 - \left(1 - \frac{1}{2^j}\right) \|\mathbf{h}_i^0 - \mathbf{h}_{k,i}^0\|^2 - \left(\frac{1}{2} - \frac{1}{2^{j+1}}\right) \|\mathbf{W}_{\text{DEQ}}\|_F^2. \quad (30)$$

So, when $j \rightarrow \infty$,

$$\begin{aligned} -\frac{K}{2} \sum_{k=1}^K \sum_{i=1}^n \|\mathcal{G}_{k,i}^0\|^2 &= \frac{K}{2} \sum_{k=1}^K \sum_{i=1}^n \left\| \sum_{j=0}^{\infty} \mathbf{W}_{\text{DEQ}}^j (\bar{\mathbf{h}}_i^0 - \mathbf{h}_{k,i}^0) \right\|^2 \\ &\geq -K \sum_{k=1}^K \sum_{i=1}^n \left(\|\mathbf{h}_i^0 - \mathbf{h}_{k,i}^0\|^2 - \frac{1}{2} \|\mathbf{W}_{\text{DEQ}}\|_F^2 \right). \end{aligned} \quad (31)$$

Therefore, use a similar proof as a fully connected layer,

$$\begin{aligned} \tilde{\mathcal{L}} &\geq -\frac{K}{C_5} \sum_{k=1}^K \sum_{i=1}^n \|\mathbf{h}_i^0 - \mathbf{h}_{k,i}^0\|^2 - \frac{NK}{2C_5} \|\mathbf{W}_{\text{DEQ}}\|_F^2 - \frac{C_5 NK}{2} E_W \\ &= -\frac{K}{C_5} \sum_{i=1}^n \left(\frac{1}{K^2} \sum_{k=1}^K \|\mathbf{h}_{k,i}^0\|^2 - \|\bar{\mathbf{h}}_i^0\|^2 \right) - \frac{NK}{2C_5} \|\mathbf{W}_{\text{DEQ}}\|_F^2 - \frac{C_5 NK}{2} E_W \\ &\geq -\frac{NK}{C_5 E_H} - \frac{C_5 NK}{2} E_W + \frac{K^2}{C_5} \sum_{i=1}^n \|\bar{\mathbf{h}}_i^0\|^2 - \frac{NK}{2C_5} E_H. \end{aligned} \quad (32)$$

Set $C_5 = \sqrt{E_H/E_W}$, the loss bound of the loss function becomes:

$$\inf \mathcal{L}_{\text{DEQ}} = -\frac{2C_1 K}{(C_1 + C_2)(K-1)} \sqrt{E_W E_H} + C_4. \quad (33)$$

In comparison with the lower bound of the loss function (22), it is evident that the loss function of the DEQ layer is significantly lower than that of the explicit neural network. Since the models are identical, according to Remark B.2, the values of C_1 and C_2 are nearly the same. This observation highlights the relatively stronger potential of DEQ compared to Explicit Neural Networks.

Also, the satisfaction conditions for the inequalities in DEQ settings include the following:

- In Eq. (19): The first inequality becomes equality when

$$\frac{(C_1 + C_2) \mathbf{h}_{k,i}^T \mathbf{w}_k}{C_1} = \frac{\mathbf{h}_{k,i}^T \mathbf{w}_{k'}}{C_3}, \quad (34)$$

that is,

$$\mathbf{h}_{k,i}^T \mathbf{w}_k = \mathbf{h}_{k,i}^T \mathbf{w}_{k'} + \log \left(\frac{C_1(K-1)}{C_2} \right). \quad (35)$$

The second inequality is reduced to equality when $\bar{\mathbf{h}}_i - \mathbf{h}_{k,i} = -C_5 \mathbf{w}_k$. This condition is quite similar to explicit fully connected layers.

- In Eq. (29):

The first inequality:

$$\mathbf{W}_{\text{DEQ}} \mathcal{G}_{k,i}^{j-1} = \bar{\mathbf{h}}_i^0 - \mathbf{h}_{k,i}^0, \quad (36)$$

and the second inequality

$$\|\mathbf{W}_{\text{DEQ}}\|_F^2 = \|\mathcal{G}_{k,i}^{j-1}\|^2. \quad (37)$$

- In Eq. (32): When the following condition $\frac{1}{K} \sum_{k=1}^K \|\mathbf{w}_k\|^2 = E_W$ and $\|\mathbf{W}_{\text{DEQ}}\|^2 = \frac{1}{K} \sum_{k=1}^K \sum_{i=1}^{n_k} \|\mathbf{h}_{k,i}\|^2 = E_H$ holds, the inequality was reduced to equality.

To summarize, DEQs are proposed for the memory-saving properties, as the forward passes can leverage any black-box root solvers [3, 5]. However, in terms of forward inference, explicit neural networks have limited learning capacity for data representation since they involve direct expressions computed in a single pass and backward propagation. In contrast, DEQ, lacking a direct explicit form,

requires multiple rounds of parameter adjustments for learning. In each iteration, DEQ introduces input data in a sequential manner, allowing more adjustment space for learning parameters specific to the input. Therefore, to compare the two loss functions, we can derive the following theorem:

Theorem B.3. *DEQ achieves a lower bound on the loss function compared to explicit neural network under balanced datasets:*

$$\inf \mathcal{L}_{DEQ} = -2C_1 \frac{K}{K-1} \sqrt{E_W E_H} + C_2,$$

while the lower bound of loss function of explicit neural network remains:

$$\inf \mathcal{L}_{EX} = -C_1 \frac{K}{K-1} \sqrt{E_W E_H} + C_2,$$

where C_1 and C_2 are two given constants.

Under the balanced dataset, the sample distribution of each class within each batch is relatively even. Therefore, during the fixed-point iteration process, both DEQ and explicit neural network can learn the features of each class relatively well, without showing significant differences. Besides, from a numerical perspective, the penalties E_W and E_H are generally not set to very large values, especially smaller than 1, so the difference between the two lower bounds in Theorem B.3 may not be substantial. Besides, as analyzed in Remark B.2, we can set C_2 in this two equations as identical, and once the proportion of logits in the explicit neural network is greater than the DEQ, the lower bound of loss function in DEQ is lower.

C Proof under imbalanced learning

C.1 Lower bound of the loss function

Consider the loss function:

$$\mathcal{L} = \underbrace{\frac{K_A n_A}{N} \sum_{k=1}^{K_A} \sum_{i=1}^{n_A} \mathcal{L}(\mathbf{W} \mathbf{h}, \mathbf{y}_k)}_{\mathcal{L}_A} + \underbrace{\frac{K_B n_B}{N} \sum_{k=K_A+1}^{K_B} \sum_{i=1}^{n_B} \mathcal{L}(\mathbf{W} \mathbf{h}, \mathbf{y}_k)}_{\mathcal{L}_B}. \quad (38)$$

First analyze the loss in the majority class \mathcal{L}_A and introduce each term in the loss function. Suppose sample i belongs to category k , where $1 \leq k \leq K_A$, i.e., k is a majority class.

By applying Jensen's inequalities, we can derive:

$$\begin{aligned} & -\log \left(\frac{\exp(\mathbf{h}_{k,i}^T \mathbf{w}_k)}{\sum_{k'=1}^K \exp(\mathbf{h}_{k,i}^T \mathbf{w}_{k'})} \right) \\ &= -\mathbf{h}_{k,i}^T \mathbf{w}_k + \log \left(C_1 \exp \left(\frac{\mathbf{h}_{k,i}^T \mathbf{w}_k}{C_1} \right) + C_2 \sum_{k' \neq k} \exp \left(\frac{\mathbf{h}_{k,i}^T \mathbf{w}_{k'}}{C_2} \right) + C_3 \sum_{k'=K_A+1}^K \exp \left(\frac{\mathbf{h}_{k,i}^T \mathbf{w}_{k'}}{C_3} \right) \right) \\ &\geq (C_1 - 1) \mathbf{h}_{k,i}^T \mathbf{w}_k + C_2 \sum_{k' \neq k} \mathbf{h}_{k,i}^T \mathbf{w}_{k'} + C_3 \sum_{k'=K_A+1}^K \mathbf{h}_{k,i}^T \mathbf{w}_{k'} + \text{const} \\ &= C_0 C_4 \left(\frac{1}{K_A} \sum_{k'=1}^{K_A} \mathbf{h}_{k,i}^T \mathbf{w}_{k'} - \mathbf{h}_{k,i}^T \mathbf{w}_k \right) + C_0 C_5 \left(\frac{1}{K_B} \sum_{k'=K_A+1}^K \mathbf{h}_{k,i}^T \mathbf{w}_{k'} - \mathbf{h}_{k,i}^T \mathbf{w}_k \right) + \text{const} \\ &= C_0 C_4 (\mathbf{h}_{k,i}^T \mathbf{w}_A - \mathbf{h}_{k,i}^T \mathbf{w}_k) + C_0 C_5 (\mathbf{h}_{k,i}^T \mathbf{w}_B - \mathbf{h}_{k,i}^T \mathbf{w}_k) + \text{const}. \end{aligned} \quad (39)$$

Here the value of const is $-C_1 \log C_1 - (k_A - 1) C_2 \log C_2 - K_B C_3 \log C_3$. Besides, $\mathbf{w}_A = \frac{1}{K_A} \sum_{k'=1}^{K_A} \mathbf{w}_{k'}$ and $\mathbf{w}_B = \frac{1}{K_B} \sum_{k'=K_A+1}^K \mathbf{w}_{k'}$ represent the mean values of the weights in majority and minority classes, respectively.

To ensure the equality conditions hold, suppose there are three adaptive constants $a > 0$, $b > 0$, $c > 0$. Denote $C_1 = \frac{a}{a + (K_A - 1)b + K_{BC}}$, $C_2 = \frac{b}{a + (K_A - 1)b + K_{BC}}$, and $C_3 = \frac{c}{a + (K_A - 1)b + K_{BC}}$. Additionally, to ensure $C_4 + C_5 = 1$, introduce a constant $C_0 = \frac{K_A b + K_{BC}}{a + (K_A - 1)b + K_{BC}}$, thus $C_4 = \frac{K_A b}{K_A b + K_{BC}}$ and $C_5 = \frac{K_{BC}}{K_A b + K_{BC}}$.

After aggregating each term in the loss function, we obtain:

$$\begin{aligned} & \frac{1}{K_A n_A} \sum_{k=1}^{K_A} \sum_{i=1}^{n_A} \mathcal{L}(\mathbf{W}\mathbf{h}, \mathbf{y}_k) \\ & \geq \frac{1}{K_A n_A} \sum_{k=1}^{K_A} \sum_{i=1}^{n_A} C_4 \left(\mathbf{h}_{k,i}^T \mathbf{w}_A - \mathbf{h}_{k,i}^T \mathbf{w}_k \right) + C_5 \left(\mathbf{h}_{k,i}^T \mathbf{w}_B - \mathbf{h}_{k,i}^T \mathbf{w}_k \right) + \text{const} \quad (40) \\ & = \frac{1}{K_A} \sum_{k=1}^{K_A} \mathbf{h}_k^T (C_4 \mathbf{w}_A + C_5 \mathbf{w}_B - \mathbf{w}_k) + \text{const}, \end{aligned}$$

where $\mathbf{h}_k = \frac{1}{n_A} \sum_{i=1}^{n_A} \mathbf{h}_{k,i}$.

Subsequently, consider the lower bound of

$$\sum_{k=1}^{K_A} \mathbf{h}_k^T (C_4 \mathbf{w}_A + C_5 \mathbf{w}_B - \mathbf{w}_k) \geq -\frac{C_6}{2} \sum_{k=1}^{K_A} \|\mathbf{h}_k\|^2 - \sum_{k=1}^{K_A} \frac{1}{2} \|C_4 \mathbf{w}_A + C_5 \mathbf{w}_B - \mathbf{w}_k\|^2 / C_6. \quad (41)$$

Note that this inequality (41) is reduced to equality only when the following equality holds:

$$C_4 \mathbf{w}_A + C_5 \mathbf{w}_B - \mathbf{w}_k = C_6 \mathbf{h}_k, \quad (42)$$

where $1 \leq k \leq K_A$.

Continuing the analysis of inequality (41), the first term on the right-hand side can be bounded as:

Case 1: (Explicit fully connected layers)

$$\begin{aligned} -\sum_{k=1}^{K_A} \|\mathbf{h}_k\|^2 &= -\sum_{k=1}^{K_A} \|\mathbf{W}_{\text{EX}} \mathbf{h}_k^0\|^2 \\ &\geq -\frac{1}{2} \left(K_A \|\mathbf{W}_{\text{EX}}\|_F + \sum_{k=1}^{K_A} \|\mathbf{h}_k^0\|^2 \right) \quad (43) \\ &\geq -\frac{1}{2} \left(K_A \|\mathbf{W}_{\text{EX}}\|_F + \sum_{k=1}^{K_A} \frac{1}{n_k} \sum_{i=1}^{n_k} \|\mathbf{h}_{k,i}^0\|^2 \right) \\ &\geq -K_A E_H. \end{aligned}$$

Case 2: (Deep Equilibrium Models)

$$\begin{aligned} -\sum_{k=1}^{K_A} \|\mathbf{h}_k\|^2 &= -\sum_{k=1}^{K_A} \|(\mathbf{I} - \mathbf{W}_{\text{DEQ}})^{-1} \mathbf{h}_k^0\|^2 \\ &\geq -\frac{1}{2} \left(K_A \sum_{j=0}^{\infty} \|\mathbf{W}_{\text{DEQ}}^j\|_F + \sum_{k=1}^{K_A} \|\mathbf{h}_k^0\|^2 \right) \quad (44) \\ &\geq -\frac{1}{2} \left(K_A \sum_{j=0}^{\infty} E_H^j + \sum_{k=1}^{K_A} \frac{1}{n_k} \sum_{i=1}^{n_k} \|\mathbf{h}_{k,i}^0\|^2 \right) \\ &\geq -\frac{1}{2} \left(\frac{1}{1 - E_H} + E_H \right). \end{aligned}$$

Compared the lower bound of explicit neural network and DEQ, we can find that:

$$\left(-\sum_{k=1}^{K_A} \|\mathbf{h}_k\|^2\right)_{\text{DEQ}} < \left(-\sum_{k=1}^{K_A} \|\mathbf{h}_k\|^2\right)_{\text{EX}}$$

for all $E_H \neq 1$.

We now shift our attention to the second term (Ref. Eq [82-83] in [14]):

$$\begin{aligned} & -\frac{1}{K_A} \sum_{k=1}^{K_A} \|C_4 \mathbf{w}_A + C_5 \mathbf{w}_B - \mathbf{w}_k\|^2 \\ = & -\frac{1}{K_A} \sum_{k=1}^{K_A} \|\mathbf{w}_k\|^2 + \frac{2}{K_A} \sum_{k=1}^{K_A} \mathbf{w}_k^T (C_4 \mathbf{w}_A + C_5 \mathbf{w}_B) - \|C_4 \mathbf{w}_A + C_5 \mathbf{w}_B\|^2 \\ = & -\frac{1}{K_A} \sum_{k=1}^{K_A} \|\mathbf{w}_k\|^2 + 2C_5^2 \mathbf{w}_A^T \mathbf{w}_B + C_4(2 - C_4) \|\mathbf{w}_A\|^2 - C_5 \|\mathbf{w}_B\|^2 \\ = & -\frac{1}{K_A} \sum_{k=1}^{K_A} \|\mathbf{w}_k\|^2 + \frac{1}{K_A} \sum_{k=K_A+1}^K \|\mathbf{w}_k\|^2 + C_4(2 - C_4) \left\| \mathbf{w}_A + \frac{C_5^2}{C_4(2 - C_4)} \mathbf{w}_B \right\|^2 \\ & - \left(C_5^2 + \frac{C_5^2}{C_4(2 - C_4)} \right) \|\mathbf{w}_B\|^2 \\ \geq & -\frac{K}{K_A} E_W + \left(\frac{1}{K_R} - C_5^2 - \frac{C_5^4}{C_4(2 - C_4)} \right) \|\mathbf{w}_B\|^2 + \frac{1}{K_A} \sum_{k=K_A+1}^K \|\mathbf{w}_k - \mathbf{w}_B\|^2 \\ & + C_4(2 - C_4) \left\| \mathbf{w}_A + \frac{C_5^2}{C_4(2 - C_4)} \mathbf{w}_B \right\|^2, \end{aligned} \quad (45)$$

where $K_R = K_A/K_B$ denotes the ratio of the number of majority classes to minority classes.

In summary, the lower bound of loss function (40) could be simplified as:

$$\begin{aligned} \mathcal{L}_A &= \frac{1}{K_A n_A} \sum_{k=1}^{K_A} \sum_{i=1}^{n_A} \mathcal{L}(\mathbf{W} \mathbf{h}_{k,i}, \mathbf{y}_k) \\ &\geq \frac{1}{K_A} \sum_{k=1}^{K_A} \mathbf{h}_k^T (C_4 \mathbf{w}_A + C_5 \mathbf{w}_B - \mathbf{w}_k) + \text{const} \\ &\geq -\frac{C_6}{2K_A} \sum_{k=1}^{K_A} \|\mathbf{h}_k\|^2 - \frac{1}{2K_A} \sum_{k=1}^{K_A} \|C_4 \mathbf{w}_A + C_5 \mathbf{w}_B - \mathbf{w}_k\|^2 / C_6 + \text{const} \\ &\geq \frac{C_6}{2K_A} M - \frac{K E_W}{2C_6 K_A} + \frac{1}{2C_6} \left(\frac{1}{K_R} - C_5^2 - \frac{C_5^4}{C_4(2 - C_4)} \right) \|\mathbf{w}_B\|^2 \\ &\quad + \frac{C_4(2 - C_4)}{C_6} \left\| \mathbf{w}_A + \frac{C_5^2}{C_4(2 - C_4)} \mathbf{w}_B \right\|^2 + \frac{1}{2C_6 K_A} \sum_{k=K_A+1}^K \|\mathbf{w}_k - \mathbf{w}_B\|^2 + \text{const}, \end{aligned} \quad (46)$$

where $M = -K_A E_H$ if the network is a fully connected layer and $M = -\frac{K_A}{2} \left(\frac{1}{1 - E_H} + E_H \right)$ if the network is a Deep Equilibrium Model.

Similarly, the loss function w.r.t the minority classes is bounded as:

$$\begin{aligned}
\mathcal{L}_B &= \frac{1}{K_B n_B} \sum_{k=1}^{K_B} \sum_{i=1}^{n_B} \mathcal{L}(\mathbf{W} \mathbf{h}_{k,i}, \mathbf{y}_k) \\
&= \frac{1}{K_B} \sum_{k=1}^{K_B} \mathbf{h}_k^T (C_4 \mathbf{w}_A + C_5 \mathbf{w}_B - \mathbf{w}_k) + \text{const} \\
&\geq -\frac{C_6}{2K_B} \sum_{k=1}^{K_B} \|\mathbf{h}_k\|^2 - \frac{1}{2K_B} \sum_{k=1}^{K_B} \|C_4 \mathbf{w}_A + C_5 \mathbf{w}_B - \mathbf{w}_k\|^2 / C_6 + \text{const} \\
&\geq \frac{C_6}{2K_B} M - \frac{KEW}{2C_6 K_B} + \frac{1}{2C_6} \left(K_R - C_5^2 - \frac{C_5^4}{C_4(2-C_4)} \right) \|\mathbf{w}_A\|^2 \\
&\quad + \frac{C_5(2-C_5)}{C_6} \left\| \frac{C_4^2}{C_5(2-C_5)} \mathbf{w}_A + \mathbf{w}_B \right\|^2 + \frac{1}{2C_6 K_B} \sum_{k=1}^{K_A} \|\mathbf{w}_k - \mathbf{w}_A\|^2 + \text{const}.
\end{aligned} \tag{47}$$

The inequality reduces to equality when the constraints in \mathcal{C} are treated as equalities, achieving the upper bound. Additionally, the following equalities must hold:

$$C_4 \mathbf{w}_A + C_5 \mathbf{w}_B - \mathbf{w}_k = C_6 \mathbf{h}_k, \tag{48}$$

where $K_A + 1 \leq k \leq K$.

If $K_R = 1$, i.e., the number of majority classes is equal to the number of minority classes, the results of (46) and (47) are totally equivalent.

Therefore, without loss of generality, assuming $K_A > K_B$, the lower bound of the loss function (38) can be simplified to:

$$\begin{aligned}
\mathcal{L} &= \mathcal{L}_A + \mathcal{L}_B \\
&\geq \frac{C_6 M}{2} \left(\frac{1}{K_A} + \frac{1}{K_B} \right) + \frac{1}{2C_6 K_B} \sum_{k=1}^{K_A} \|\mathbf{w}_k - \mathbf{w}_A\|^2 + \frac{1}{2C_6 K_A} \sum_{k=1}^{K_B} \|\mathbf{w}_k - \mathbf{w}_B\|^2 \\
&\quad + \frac{C_4(2-C_4)}{2C_6} \left\| \mathbf{w}_A + \frac{C_5^2}{C_4(2-C_4)} \mathbf{w}_B \right\|^2 + \frac{C_5(2-C_5)}{2C_6} \left\| \frac{C_4^2}{C_5(2-C_5)} \mathbf{w}_A + \mathbf{w}_B \right\|^2 \\
&\quad + \frac{1}{2C_6} \left(\frac{1}{K_R} - C_5^2 - \frac{C_5^4}{C_4(2-C_4)} \right) \|\mathbf{w}_B\|^2 + \frac{1}{2C_6} \left(K_R - C_5^2 - \frac{C_5^4}{C_4(2-C_4)} \right) \|\mathbf{w}_A\|^2 + \text{const}.
\end{aligned} \tag{49}$$

C.2 \mathcal{NC} Analysis

As analyzed in (43) and (44), when it reaches the minimal value, each $\mathbf{h}_{k,i} = \mathbf{h}_k$ for $\forall k = 1, 2, \dots, K_A$. Similarly, this holds for minority class with $K_A + 1 \leq k \leq K$. This implies that, in an imbalanced scenario, both DEQ and fully connected layer exhibit feature collapse, i.e., $\mathcal{NC}1$ is still present.

As we need to calculate the lower bound of the loss function, it is essential to minimize the terms $\mathcal{L}_A + \mathcal{L}_B$ as much as possible.

Therefore, consider the gradient with respect to \mathbf{w}_k for majority class and \mathbf{w}_k for minority class, respectively. First compute the case with $1 \leq k \leq K_A$.

$$\begin{aligned}
\frac{\partial \mathcal{L}}{\partial \mathbf{w}_k} &= \frac{1}{C_6 K_B} \left(1 - \frac{1}{K_A} \right) (\mathbf{w}_k - \mathbf{w}_A) + \frac{C_4(2-C_4)}{C_6 K_A} \left(\mathbf{w}_A + \frac{C_5^2}{C_4(2-C_4)} \mathbf{w}_B \right) \\
&\quad + \frac{C_4^2}{K_A C_6} \left(\frac{C_4^2}{C_5(2-C_5)} \mathbf{w}_A + \mathbf{w}_B \right) \\
&\quad + \frac{1}{K_A C_6} \left(K_R - C_5^2 - \frac{C_5^4}{C_4(2-C_4)} \right) \mathbf{w}_A = 0.
\end{aligned} \tag{50}$$

So, we can derive that

$$\begin{aligned} \left(K_R - \frac{1}{K_B}\right) \mathbf{w}_k + \frac{1}{1 - C_4^2} \mathbf{w}_B \\ + \left(\frac{1}{K_B} + C_4(2 - C_4) + \frac{C_4^4}{C_5(2 - C_5)} - C_5^2 - \frac{C_5^2}{C_4(2 - C_4)}\right) \mathbf{w}_A = 0. \end{aligned} \quad (51)$$

One important note here is that when the proportion of majority class samples approaches infinity, i.e., $C_4 \rightarrow 1$, we have $\frac{1}{1 - C_4^2} \rightarrow 0$. In this scenario, the weights \mathbf{w}_k belonging to the majority class are almost exclusively related to \mathbf{w}_A , and have little dependence on the average of the minority class \mathbf{w}_B , which validates the results of Proposition 4.2.

Similarly, if $K_A + 1 \leq k \leq K$, the following equality will hold to ensure optimality of \mathbf{w}_k in minority classes of (49):

$$\begin{aligned} \left(\frac{1}{K_R} - \frac{1}{K_A}\right) \mathbf{w}_k + \frac{1}{1 - C_5^2} \mathbf{w}_A \\ + \left(\frac{1}{K_A} + C_5(2 - C_5) + \frac{C_5^4}{C_4(2 - C_4)} - C_4^2 - \frac{C_4^2}{C_5(2 - C_5)}\right) \mathbf{w}_B = 0. \end{aligned} \quad (52)$$

Next, we consider the conditions for the validity in $\mathcal{NC}2$ and $\mathcal{NC}3$, then compare the performance of DEQ and explicit neural network.

Therefore, for the majority class $1 \leq k \leq K_A$, suppose it reaches its minimum value, recall the condition (42), and combined with (51), we can derive:

$$C_6 \mathbf{h}_k^T \mathbf{h}_{k'} = \left(C_4 + \frac{K_B C_A}{K_A - 1}\right) \mathbf{w}_A^T \mathbf{h}_{k'} + \left(C_5 + \frac{K_B}{(K_A - 1)(1 - C_4^2)}\right) \mathbf{w}_B^T \mathbf{h}_{k'}. \quad (53)$$

Similarly, for the minority class $K_A + 1 \leq k \leq K$, combine (48) with (52), we obtain:

$$C_6 \mathbf{h}_k^T \mathbf{h}_{k'} = \left(C_4 + \frac{K_A}{(K_B - 1)(1 - C_5^2)}\right) \mathbf{w}_A^T \mathbf{h}_{k'} + \left(C_5 + \frac{K_A C_B}{K_B - 1}\right) \mathbf{w}_B^T \mathbf{h}_{k'}. \quad (54)$$

In the above two equations, $k' = 1, 2, \dots, K$. And we denote the coefficient of \mathbf{w}_A in Eq. (51) and the coefficient of \mathbf{w}_B in Eq. (52) as C_A and C_B respectively for simplicity. After this deviation, we can find that both of the coefficients of $\mathbf{w}_A^T \mathbf{h}_{k'}$ and $\mathbf{w}_B^T \mathbf{h}_{k'}$ are constants.

It can be obviously concluded that $\mathcal{NC}2$ and $\mathcal{NC}3$ do not hold under imbalanced dataset conditions. However, we can still compare the numerical differences between them under DEQ and fully connected layer. By adaptively specifying parameters C_4 and C_5 , we can denote $(\mathbf{h}_k^0)^T \mathbf{h}_{k'}^0 = \mathbf{m}_{k,k'}$.

Thus, by considering all the equality conditions in (46), we can measure the distance from the features to the Simplex ETF.

$$C_6 \mathbf{h}_{k'}^T \mathbf{h}_k = C_4 \mathbf{h}_{k'}^T \mathbf{w}_A + C_5 \mathbf{h}_{k'}^T \mathbf{w}_B - \mathbf{h}_{k'}^T \mathbf{w}_k. \quad (55)$$

Case 1: (Explicit fully connected layers)

$$\begin{aligned} C_6 (\mathbf{h}_{k'}^0)^T \mathbf{h}_k &= C_4 \mathbf{W}_{\text{EX}} (\mathbf{h}_{k'}^0)^T \mathbf{w}_A + C_5 \mathbf{W}_{\text{EX}} (\mathbf{h}_{k'}^0)^T \mathbf{w}_B - \mathbf{W}_{\text{EX}} (\mathbf{h}_{k'}^0)^T \mathbf{w}_k \\ &= \mathbf{W}_{\text{EX}} \left(C_4 \mathbf{h}_{k'}^T \mathbf{w}_A + C_5 \mathbf{h}_{k'}^T \mathbf{w}_B - \mathbf{h}_{k'}^T \mathbf{w}_k \right) \\ &\leq \frac{1}{2} \|\mathbf{W}_{\text{EX}}\|_F + \frac{1}{2} \left\| C_4 \mathbf{h}_{k'}^T \mathbf{w}_A + C_5 \mathbf{h}_{k'}^T \mathbf{w}_B - \mathbf{h}_{k'}^T \mathbf{w}_k \right\| \\ &= E_H + \frac{1}{2} \mathbf{M}. \end{aligned} \quad (56)$$

Case 2: (Deep Equilibrium Models)

$$\begin{aligned}
C_6(h_{k'}^0)^T \mathbf{h}_k &= C_4(\mathbf{I} - \mathbf{W}_{\text{DEQ}})^{-1}(h_{k'}^0)^T \mathbf{w}_A + C_5(\mathbf{I} - \mathbf{W}_{\text{DEQ}})^{-1}(h_{k'}^0)^T \mathbf{w}_B \\
&\quad - (\mathbf{I} - \mathbf{W}_{\text{DEQ}})^{-1}(h_{k'}^0)^T \mathbf{w}_k \\
&= (\mathbf{I} - \mathbf{W}_{\text{DEQ}})^{-1} \left(C_4 \mathbf{h}_{k'}^T \mathbf{w}_A + C_5 \mathbf{h}_{k'}^T \mathbf{w}_B - \mathbf{h}_{k'}^T \mathbf{w}_k \right) \\
&\leq \frac{1}{2} \|(\mathbf{I} - \mathbf{W}_{\text{DEQ}})^{-1}\|_F + \frac{1}{2} \left\| C_4 \mathbf{h}_{k'}^T \mathbf{w}_A + C_5 \mathbf{h}_{k'}^T \mathbf{w}_B - \mathbf{h}_{k'}^T \mathbf{w}_k \right\| \\
&= \frac{1}{2(1 - E_H)} + \frac{1}{2} \mathbf{M}.
\end{aligned} \tag{57}$$

Therefore, to compare these two models, we consider the case when the distance of these two models from the Simplex ETF is minimized. We denote each element in the Simplex ETF as s_{ij} and compare the differences between them. When the distance of DEQ is relatively smaller than that of explicit neural network, we can obtain:

$$\left| \frac{1}{2(1 - E_H)} + \frac{1}{2} \mathbf{m} - \mathbf{s} \right| < \left| \frac{1}{2} E_H + \frac{1}{2} \mathbf{m} - \mathbf{s} \right|. \tag{58}$$

For simplicity, we only consider the subscripts of \mathbf{s} and \mathbf{m} , and denote ① = $\frac{1}{2(1 - E_H)} + \frac{1}{2} \mathbf{m} - \mathbf{s}$ and ② = $\frac{1}{2} E_H + \frac{1}{2} \mathbf{m} - \mathbf{s}$. We then classify and discuss their magnitudes.

- ① > 0, ② > 0:
Since $\frac{1}{2(1 - E_H)} > E_H$, there is a contradiction! Therefore, it does not hold.
- ① > 0, ② < 0:
 $\frac{1}{2(1 - E_H)} + \frac{1}{2} \mathbf{m} - \mathbf{s} < 0$, $\frac{1}{2} E_H + \frac{1}{2} \mathbf{m} - \mathbf{s} > 0$, which means $\frac{1}{2(1 - E_H)} < E_H$. And that is a contradiction!
- ① < 0, ② > 0:
Since $\frac{1}{2(1 - E_H)} + \frac{1}{2} \mathbf{m} - \mathbf{s} > 0$, $\frac{1}{2} E_H + \frac{1}{2} \mathbf{m} - \mathbf{s} < 0$, we have

$$E_H < 2\mathbf{s} - \mathbf{m} < \frac{1}{1 - E_H}.$$

Besides, by the inequality (58), we have $\frac{1}{2} E_H + \frac{1}{2(1 - E_H)} < 2\mathbf{s} - \mathbf{m}$. Then find the intersection, we obtain that:

$$\frac{1}{2} E_H + \frac{1}{2(1 - E_H)} < 2\mathbf{s} - \mathbf{m} < \frac{1}{1 - E_H}.$$

- ① < 0, ② < 0:
Since $\frac{1}{2(1 - E_H)} + \frac{1}{2} \mathbf{m} - \mathbf{s} < 0$ and $\frac{1}{2} E_H + \frac{1}{2} \mathbf{m} - \mathbf{s} < 0$, it implies that $\frac{1}{2(1 - E_H)} > E_H$ always holds.

Therefore, we only need to ensure that

$$2\mathbf{s} - \mathbf{m} > \min \left\{ E_H, \frac{1}{1 - E_H} \right\}.$$

Combining these four cases and finding their intersection, we conclude that when the inequality

$$E_H < 2\mathbf{s} - \mathbf{m} < \frac{1}{1 - E_H} \tag{59}$$

is satisfied, the performance of DEQ is better than that of explicit neural network.

As for $\mathcal{NC3}$, consider the cosine distance with the feature \mathbf{h}_k and \mathbf{w}_k .

Case 1: (Explicit fully connected layers)

$$\begin{aligned}
 \cos(\mathbf{h}_k, \mathbf{w}_k)_{\text{EX}} &= \frac{\mathbf{h}_k^T \mathbf{w}_k}{\|\mathbf{w}_k\| \|\mathbf{h}_k\|} \\
 &= \frac{\mathbf{W}_{\text{EX}} (\mathbf{h}_k^0)^T \mathbf{w}_k}{\|\mathbf{w}_k\| \|\mathbf{W}_{\text{EX}} \mathbf{h}_k^0\|} \\
 &\geq \frac{2\mathbf{W}_{\text{EX}} (\mathbf{h}_k^0)^T \mathbf{w}_k}{\|\mathbf{w}_k\|^2 + \frac{1}{2}\|\mathbf{W}_{\text{EX}}\|^2 + \frac{1}{2}\|\mathbf{h}_k^0\|^2} \\
 &\geq \frac{2E_H (\mathbf{h}_k^0)^T \mathbf{w}_k}{E_W + E_H}.
 \end{aligned} \tag{60}$$

Case 2: (Deep Equilibrium Model)

Very similarly,

$$\begin{aligned}
 \cos(\mathbf{h}_k, \mathbf{w}_k)_{\text{DEQ}} &= \frac{\mathbf{h}_k^T \mathbf{w}_k}{\|\mathbf{w}_k\| \|\mathbf{h}_k\|} \\
 &= \frac{(\mathbf{I} - \mathbf{W}_{\text{DEQ}}) (\mathbf{h}_k^0)^T \mathbf{w}_k}{\|\mathbf{w}_k\| \|(\mathbf{I} - \mathbf{W}_{\text{DEQ}})^{-1} \mathbf{h}_k^0\|} \\
 &\geq \frac{2(\mathbf{I} - \mathbf{W}_{\text{DEQ}}) (\mathbf{h}_k^0)^T \mathbf{w}_k}{\|\mathbf{w}_k\|^2 + \frac{1}{2}\|(\mathbf{I} - \mathbf{W}_{\text{DEQ}})\|^2 + \frac{1}{2}\|\mathbf{h}_k^0\|^2} \\
 &\geq \frac{4E_H (\mathbf{h}_k^0)^T \mathbf{w}_k}{1 + 2(E_W + E_H)(1 - E_H)}.
 \end{aligned} \tag{61}$$

If the performance of DEQ is better than explicit neural network, then we have

$$\cos(\mathbf{h}_k, \mathbf{w}_k)_{\text{DEQ}} / \cos(\mathbf{h}_k, \mathbf{w}_k)_{\text{exp}} > 1,$$

which is equivalent to

$$\frac{E_H}{E_W + E_H} + E_H(1 - E_H) < 2. \tag{62}$$

In summary, though DEQ does not completely mitigate the issue of minority collapse, it shows significant improvement compared to explicit neural network under some conditions that are relatively easy to satisfy in the manifestation of the \mathcal{NC} phenomenon.

D More experiments

In this section, we provide more experimental results, including the \mathcal{NC} phenomena of Explicit NN and DEQ, and the training results under other imbalanced conditions.

Table 3: Test Accuracy on Cifar-10 and Cifar-100 Dataset with $K_A = 5$

	R	Cifar-10			Cifar-100		
		10	50	100	10	50	100
Explicit NN	overall	80.73±0.48	63.08±0.87	44.86±1.43	52.62±0.86	41.62±0.68	37.33±2.29
	majority	94.18±0.56	91.02±0.89	89.32±0.79	74.10±1.03	73.94±0.25	74.24±1.13
	minority	67.80±0.35	35.14±0.65	0.40±3.86	31.10±0.70	9.30±1.10	0.42±3.04
DEQ	overall	81.36±1.03	65.03±1.90	46.09±1.77	53.31±0.98	44.07±2.04	39.11±2.46
	majority	93.14±1.81	90.88±2.83	90.20±0.85	72.90±1.65	75.98±1.75	75.79±0.96
	minority	69.58±0.66	39.18±1.46	1.26±4.93	33.72±0.79	12.16±3.75	2.42±5.89

Table 4: Test Accuracy on Cifar-10 and Cifar-100 Dataset with $K_A = 7$

	R	Cifar-10			Cifar-100		
		10	50	100	10	50	100
Explicit NN	overall	83.17±0.40	66.91±0.39	53.27±0.81	59.11±0.84	51.71±1.02	50.72±0.60
	majority	89.09±0.36	80.90±0.57	75.12±0.74	71.93±0.65	72.20±0.66	72.46±0.58
	minority	69.37±0.49	34.26±0.30	2.30±1.01	29.20±0.92	3.90±1.29	0.00±0.00
DEQ	overall	83.78±1.85	69.47±1.86	56.74±0.98	60.51±0.88	52.99±1.86	51.79±0.92
	majority	88.98±1.99	82.91±2.22	78.81±0.67	72.90±1.19	72.99±0.98	73.98±0.66
	minority	71.65±1.63	38.12±1.61	5.20±1.91	31.13±0.83	6.33±2.35	0.00±0.00

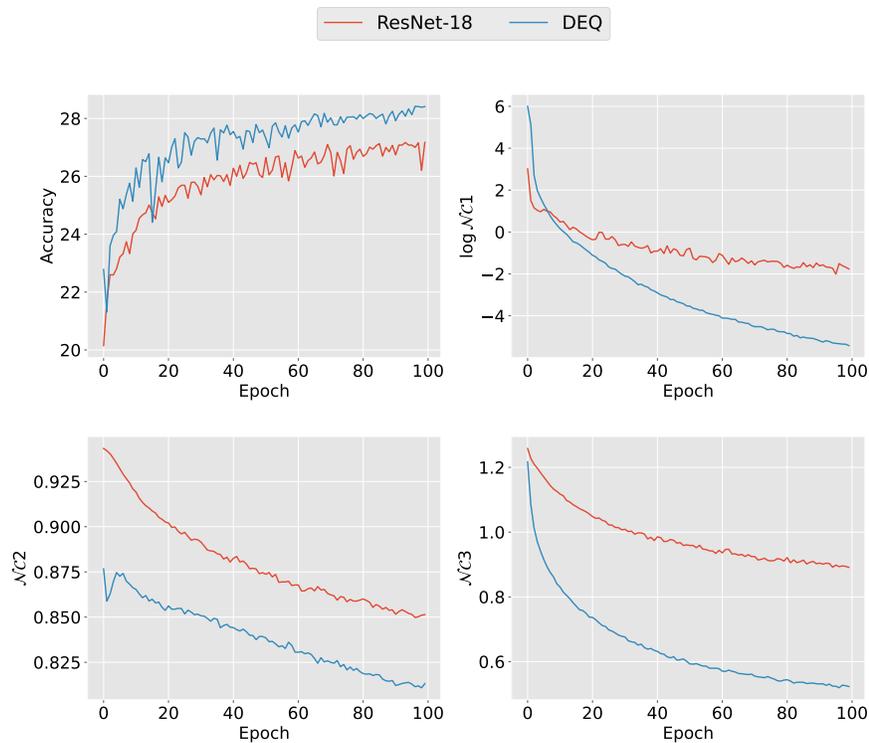


Figure 4: Accuracy and \mathcal{N}^C phenomenon on imbalanced dataset with $K_A = 3$, $K_B = 7$, $R = 100$

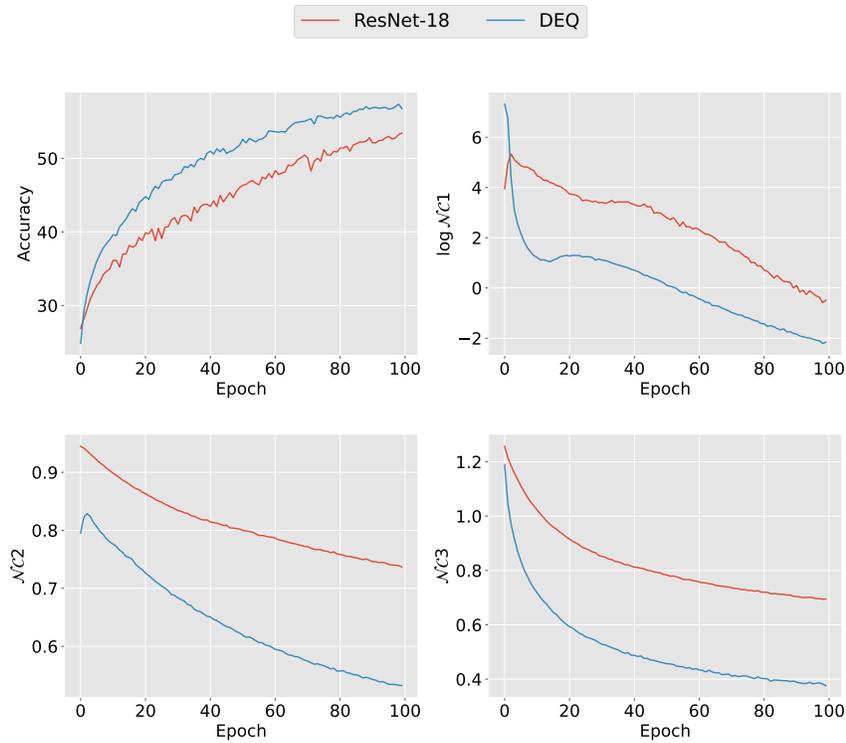


Figure 5: Accuracy and $\mathcal{N}\mathcal{C}$ phenomenon on imbalanced dataset with $K_A = 7$, $K_B = 3$, $R = 100$

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: As stated in the abstract and introduction, this paper is the first to analyze the representation of the Deep Equilibrium Model from the perspective of Neural Collapse, accurately reflecting the key contributions and scope.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The limitation analysis is provided in the Conclusion.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: The assumptions and proofs are provided in appendix B and C.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: These details are provided in Section 5.1 - Experiment setup.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We will release the code once the paper is accepted.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: These details are provided in Section 5.1 - Experiment setup.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: The standard deviation in our experimental results (Table 1-4) shows the statistical significance.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: The computer resources are provided in Section 5.1 - Experiment setup.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: Our research conformed with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: Our paper primarily focuses on theoretical research in machine learning, comparing two typical neural network algorithms, with no relevance to societal impacts.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our paper uses standard CIFAR-10 and CIFAR-100 datasets, which do not involve such issues.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: The datasets CIFAR-10 and CIFAR-100 are cited properly. Other assets are not applied in this paper.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.

- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: This paper does not introduce new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Our paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Our paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.