

On the Equivalence between Implicit and Explicit Neural Networks: A High-dimensional Viewpoint

Zenan Ling
Zhenyu Liao
Robert C. Qiu

Huazhong University of Science and Technology, Wuhan, China

LINGZENAN@HUST.EDU.CN
ZHENYU_LIAO@HUST.EDU.CN
CAIMING@HUST.EDU.CN

Abstract

Implicit neural networks have demonstrated remarkable success in various tasks. However, there is a lack of theoretical analysis of the connections and differences between implicit and explicit networks. In this paper, we study high-dimensional implicit neural networks and provide the high dimensional equivalents for the corresponding conjugate kernels and neural tangent kernels. Built upon this, we establish the equivalence between implicit and explicit networks in high dimensions.

1. Introduction

Implicit neural networks (NNs) [2] have recently emerged as a new paradigm in neural network design. An implicit NN is equivalent to an infinite-depth weight-shared explicit NN with input-injection. Unlike explicit NNs, implicit NNs generate features by directly solving for the fixed point, rather than through layer-by-layer forward propagation. Moreover, implicit NNs have the remarkable advantage that gradients can be computed analytically only through the fixed point with *implicit differentiation*. Therefore, training implicit NNs only requires constant memory.

Despite the empirical success achieved by implicit NNs [3, 11], our theoretical understanding of these models is still limited. In particular, there is a lack of theoretical analysis of the training dynamics and generalization performance of implicit NNs, and possibly more importantly, whether these properties can be connected to those of explicit NNs. [2] demonstrates that any deep NN can be reformulated as a special implicit NN. However, it remains unknown whether general implicit NNs have advantages over explicit NNs. [6] extends previous neural tangent kernel (NTK) studies to implicit NNs and give the exact expression of the NTK of the ReLU implicit NNs. However, the differences between implicit and explicit NTKs are not analyzed. Moreover, previous works [9, 10] have proved the global convergence of gradient descent for training implicit NNs. However, it is still unclear what distinguishes the training dynamic of implicit NNs and that of explicit NNs.

In this paper, we investigate implicit NNs from a high-dimensional view. Specifically, we perform a fine-grained asymptotic analysis on the eigenspectra of conjugate kernel (CKs) and NTKs of implicit NNs, which play a fundamental role in the convergence and generalization high dimensional NNs [8]. By considering input data uniformly drawn from the unit sphere, we derive, with recent advances in random matrix theory, high-dimensional (spectral) equivalents for the CKs and NTKs of implicit NNs, and establish the equivalence between implicit and explicit NNs by matching the coefficients of the corresponding asymptotic spectral equivalents. Surprisingly, our results reveal that a *single-layer* explicit NN with carefully designed activations has the same CK or NTK eigenspectra as a ReLU implicit NN, whose depth is essentially *infinite*.

2. Preliminaries

2.1. Implicit and Explicit NNs

Implicit NNs. In this paper, we study a typical implicit neural network, the deep equilibrium model (DEQ) [2]. Let $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n] \in \mathbb{R}^{d \times n}$ denote the input data. We define a vanilla DEQ with the transform at the l -th layer as

$$\mathbf{h}_i^{(l)} = \sqrt{\frac{\sigma_a^2}{m}} \mathbf{A} \mathbf{z}_i^{(l-1)} + \sqrt{\frac{\sigma_b^2}{m}} \mathbf{B} \mathbf{x}_i, \quad \mathbf{z}_i^{(l)} = \phi(\mathbf{h}_i^{(l)}) \quad (1)$$

where $\mathbf{A} \in \mathbb{R}^{m \times m}$ and $\mathbf{B} \in \mathbb{R}^{m \times d}$ are weight matrices, $\sigma_a, \sigma_b \in \mathbb{R}$ are constants, ϕ is an element-wise activation, $\mathbf{h}_i^{(l)}$ is the pre-activation and $\mathbf{z}_i^{(l)} \in \mathbb{R}^m$ is the output feature of the l -th hidden layer corresponding to the input data \mathbf{x}_i . The output of the last hidden layer is defined by $\mathbf{z}_i^* \triangleq \lim_{l \rightarrow \infty} \mathbf{z}_i^{(l)}$ and we denote the corresponding pre-activation by \mathbf{h}_i^* . Note that \mathbf{z}_i^* can be calculated by directly solving for the equilibrium point of the following equation

$$\mathbf{z}_i^* = \phi \left(\sqrt{\frac{\sigma_a^2}{m}} \mathbf{A} \mathbf{z}_i^* + \sqrt{\frac{\sigma_b^2}{m}} \mathbf{B} \mathbf{x}_i \right). \quad (2)$$

We are interested in the conjugate kernel and neural tangent kernel (Implicit-CK and Implicit-NTK, for short) of implicit neural networks defined in Eq. (2). Following [6], we denote the corresponding Implicit-CK by $\mathbf{G}^* = \lim_{l \rightarrow \infty} \mathbf{G}^{(l)}$ where the (i, j) -th entry of $\mathbf{G}^{(l)}$ is defined recursively as

$$\begin{aligned} \mathbf{G}_{ij}^{(0)} &= \mathbf{x}_i^\top \mathbf{x}_j, \quad \mathbf{\Lambda}_{ij}^{(l)} = \begin{bmatrix} \mathbf{G}_{ii}^{(l-1)} & \mathbf{G}_{ij}^{(l-1)} \\ \mathbf{G}_{jj}^{(l-1)} & \mathbf{G}_{jj}^{(l-1)} \end{bmatrix}, \\ \mathbf{G}_{ij}^{(l)} &= \sigma_a^2 \mathbb{E}_{(\mathbf{u}, \mathbf{v}) \sim \mathcal{N}(0, \mathbf{\Lambda}_{ij}^{(l)})} [\phi(\mathbf{u}) \phi(\mathbf{v})] + \sigma_b^2 \mathbf{x}_i^\top \mathbf{x}_j, \\ \dot{\mathbf{G}}_{ij}^{(l)} &= \sigma_a^2 \mathbb{E}_{(\mathbf{u}, \mathbf{v}) \sim \mathcal{N}(0, \mathbf{\Lambda}_{ij}^{(l)})} [\phi'(\mathbf{u}) \phi'(\mathbf{v})]. \end{aligned} \quad (3)$$

And the Implicit-NTK is defined as $\mathbf{K}^* = \lim_{l \rightarrow \infty} \mathbf{K}^{(l)}$ whose the (i, j) -th entry is defined as

$$\mathbf{K}_{ij}^{(l)} = \sum_{h=1}^{l+1} \left(\mathbf{G}_{ij}^{(h-1)} \prod_{h'=h}^{l+1} \dot{\mathbf{G}}_{ij}^{(h')} \right). \quad (4)$$

Explicit Neural Networks. We consider a single-layer fully-connected NN model defined as $\mathbf{Y} = \sqrt{\frac{1}{p}} \sigma(\mathbf{W} \mathbf{X})$ where $\mathbf{W} \in \mathbb{R}^{p \times d}$ is the weight matrix and σ is an element-wise activation function. Let $\mathbf{w} \sim \mathcal{N}(0, \mathbf{I}_d)$, the corresponding Explicit-CK matrix $\mathbf{\Sigma}$ and Explicit-NTK matrix $\mathbf{\Theta}$ are defined as follows:

$$\mathbf{\Sigma} = \mathbb{E}_{\mathbf{w}} [\sigma(\mathbf{w}^\top \mathbf{X})^\top \sigma(\mathbf{w}^\top \mathbf{X})], \quad \mathbf{\Theta} = \mathbf{\Sigma} + \left(\mathbf{X}^\top \mathbf{X} \right) \odot \mathbb{E}_{\mathbf{w}} [\sigma'(\mathbf{w}^\top \mathbf{X})^\top \sigma'(\mathbf{w}^\top \mathbf{X})]. \quad (5)$$

2.2. CKs and NTKs of ReLU Implicit NNs

We make the following assumptions on the random initialization, the input data, and activations.

Assumption 1 (i) As $n \rightarrow \infty$, $d/n \rightarrow c \in (0, \infty)$. All data points \mathbf{x}_i , $i \in [n]$, are independent and uniformly sampled from \mathbb{S}^{d-1} . (ii) \mathbf{A} , \mathbf{B} , and \mathbf{W} are independent and have i.i.d entries of zero mean, unit variance, and finite fourth kurtosis. Moreover, we require $\sigma_a^2 + \sigma_b^2 = 1$. (iii) The activation ϕ of the implicit NN is the normalized ReLU, i.e., $\phi(x) = \sqrt{2} \max(x, 0)$. The activation σ of the explicit NN is a C^3 function.

Remark 1 (i) Despite derived here for uniform distribution on the unit sphere, we conjecture that our results extend the result to more general distributions by using the technique developed in [5, 7]. (ii) The additional requirement on the variance is to ensure the existence and uniqueness of the fixed point of the NTK and to keep the diagonal entries of the CK matrix at 1, see examples in [6]. (iii) It is possible to extend our results to implicit NNs with general activations by using the technique proposed in [10]. We defer the extension to more general data distributions and activation functions to future work.

Under Assumptions 1, the limits of Implicit-CK and Implicit-NTK exist, and one can have precise expressions of \mathbf{G}^* and \mathbf{K}^* as follows [6, 9].

Lemma 1 Let $f(x) = \frac{\sqrt{1-x^2} + (\pi - \arccos x)x}{\pi}$. Under Assumptions 1, the fixed point of Implicit-CK \mathbf{G}_{ij}^* is the root of

$$\mathbf{G}_{ij}^* = \sigma_a^2 f(\mathbf{G}_{ij}^*) + (1 - \sigma_a^2) \mathbf{x}_i^\top \mathbf{x}_j. \quad (6)$$

The limit of Implicit-NTK is

$$\mathbf{K}_{ij}^* = h(\mathbf{G}_{ij}^*) \triangleq \frac{\mathbf{G}_{ij}^*}{1 - \mathbf{G}_{ij}^*} \quad \text{where} \quad \dot{\mathbf{G}}_{ij}^* \triangleq \sigma_a^2 \pi^{-1} (\pi - \arccos(\mathbf{G}_{ij}^*)). \quad (7)$$

3. Main Results

In this section, we prove the high-dimensional equivalents for CKs and NTKs of implicit and explicit NNs. As a result, by matching the coefficients of the asymptotic spectral equivalents, we establish the equivalence between implicit and explicit NNs in high dimensions.

3.1. Asymptotic Approximations

CKs. We begin by defining several quantities that are crucial to our results. Note that the unique fixed point of Eq. (6) exists as long as $\sigma_a^2 < 1$. We define the implicit map induced from Eq. (6) as $\mathbf{G}_{ij}^* \triangleq g(\mathbf{x}_i^\top \mathbf{x}_j)$. Let $\mathcal{L}^* = g(0)$ be the solution of $\mathcal{L}^* = \sigma_a^2 f(\mathcal{L}^*)$ when $\mathbf{x}_i^\top \mathbf{x}_j = 0$. Using implicit differentiation, one can obtain that

$$g'(0) = \frac{1 - \sigma_a^2}{1 - \sigma_a^2 f'(\mathcal{L}^*)}, \quad g''(0) = \frac{\sigma_a^2 (1 - \sigma_a^2)^2 f''(\mathcal{L}^*)}{(1 - \sigma_a^2 f'(\mathcal{L}^*))^3}.$$

Now we are ready to present the asymptotic equivalent of the Implicit-CK matrix.

Theorem 1 (Asymptotic approximation of Implicit-CKs) Let Assumptions 1 hold. As $n, d \rightarrow \infty$, the Implicit-CK matrix \mathbf{G}^* defined in Eq. (6) can be approximated consistently in operator norm, by the matrix $\overline{\mathbf{G}}$, that is $\|\mathbf{G}^* - \overline{\mathbf{G}}\|_2 \rightarrow 0$, where

$$\overline{\mathbf{G}} = \alpha \mathbf{1}\mathbf{1}^\top + \beta \mathbf{X}^\top \mathbf{X} + \mu \mathbf{I}_n,$$

with $\alpha = g(0) + \frac{g''(0)}{2d}$, $\beta = g'(0)$, and $\mu = g(1) - g(0) - g'(0)$.

Theorem 2 (Asymptotic approximation for Explicit-CKs) *Let Assumptions 1 hold. As $n, d \rightarrow \infty$, the Explicit-CK matrix Σ defined in Eq. (5) can be approximated consistently in operator norm, by the matrix $\bar{\Sigma}$, that is $\|\Sigma - \bar{\Sigma}\|_2 \rightarrow 0$, where*

$$\bar{\Sigma} = \alpha_1 \mathbf{1}\mathbf{1}^\top + \beta_1 \mathbf{X}^\top \mathbf{X} + \mu_1 \mathbf{I}_n,$$

with $\alpha_1 = \mathbb{E}[\sigma(z)]^2 + \frac{\mathbb{E}[\sigma''(z)]^2}{2d}$, $\beta_1 = \mathbb{E}[\sigma'(z)]^2$, and $\mu_1 = \mathbb{E}[\sigma^2(z)] - \mathbb{E}[\sigma(z)]^2 - \mathbb{E}[\sigma'(z)]^2$, for $z \sim \mathcal{N}(0, 1)$.

NTKs. For the Implicit-NTK, we define $\mathbf{K}_{ij}^* = k(\mathbf{x}_i^\top \mathbf{x}_j)$, i.e., $k(\mathbf{x}_i^\top \mathbf{x}_j) = h(g(\mathbf{x}_i^\top \mathbf{x}_j))$, for $i, j \in [n]$. It is easy to check that $k(0) = h(\mathcal{L}^*)$ and $k(1) = h(g(1))$. Using implicit differentiation again, we have

$$k'(0) = \frac{(1 - \sigma_a^2)h'(\mathcal{L}^*)}{\sigma_a^2 f'(\mathcal{L}^*) - 1}, \quad k''(0) = \frac{(1 - \sigma_a^2)^2(h''(\mathcal{L}^*) - \sigma_a^2 f'(\mathcal{L}^*)h''(\mathcal{L}^*) + \sigma_a^2 h'(\mathcal{L}^*)f''(\mathcal{L}^*))}{(1 - \sigma_a^2 f'(\mathcal{L}^*))^3}.$$

Now we are ready to present the asymptotic equivalent of the Implicit-NTK matrix.

Theorem 3 (Asymptotic approximation for Implicit-NTKs) *Let Assumptions 1 hold. As $n, d \rightarrow \infty$, the Implicit-NTK matrix \mathbf{K}^* defined Eq. (7) in can be approximated consistently in operator norm, by the matrix $\bar{\mathbf{K}}$, that is $\|\mathbf{K}^* - \bar{\mathbf{K}}\|_2 \rightarrow 0$, where*

$$\bar{\mathbf{K}} = \dot{\alpha} \mathbf{1}\mathbf{1}^\top + \dot{\beta} \mathbf{X}^\top \mathbf{X} + \dot{\mu} \mathbf{I}_n,$$

with $\dot{\alpha} = k(0) + \frac{k''(0)}{2d}$, $\dot{\beta} = k'(0)$, and $\dot{\mu} = k(1) - k(0) - k'(0)$.

Theorem 4 (Asymptotic approximation for Explicit-NTKs) *Let Assumptions 1 hold. As $n, d \rightarrow \infty$, the Explicit-NTK matrix Θ defined in Eq. (5) can be approximated consistently in operator norm, by the matrix $\bar{\Theta}$, that is $\|\Theta^* - \bar{\Theta}\|_2 \rightarrow 0$, where*

$$\bar{\Theta} = \dot{\alpha}_1 \mathbf{1}\mathbf{1}^\top + \dot{\beta}_1 \mathbf{X}^\top \mathbf{X} + \dot{\mu}_1 \mathbf{I}_n,$$

with $\dot{\alpha}_1 = \mathbb{E}[\sigma(z)]^2 + \frac{3\mathbb{E}[\sigma''(z)]^2}{2d}$, $\dot{\beta}_1 = 2\mathbb{E}[\sigma'(z)]^2$, and $\dot{\mu}_1 = \mathbb{E}[\sigma^2(z)] + \mathbb{E}[\sigma'(z)]^2 - \mathbb{E}[\sigma(z)]^2 - 2\mathbb{E}[\sigma'(z)]^2$ for $z \sim \mathcal{N}(0, 1)$.

Remark 2 (i) *Due to the homogeneity of the ReLU function, the Implicit-CK and the Implicit-NTK are essentially inner product kernel random matrices. Consequently, Theorem 1 and 3 can be built upon the results in [4]. We postpone the study on general activations to future work. (ii) The results in Theorem 2 and 4 generalize those of [1, 7] to the cases of “non-centred” activations, i.e., we do not require $\mathbb{E}[\sigma(z)] = 0$ for $z \sim \mathcal{N}(0, 1)$.*

3.2. The Equivalence between Implicit and Explicit NNS

In the following corollary, we show a concrete case of a single-layer explicit NN with an quadratic activation, that matches the CK or NTK eigenspectra of a ReLU implicit NN. The idea is to utilize the results of Theorems 1-4 to match the coefficients of the asymptotic equivalents such that $\alpha_1 = \dot{\alpha}$, $\beta_1 = \dot{\beta}$, $\mu_1 = \dot{\mu}$, or $\dot{\alpha}_1 = \dot{\alpha}$, $\dot{\beta}_1 = \dot{\beta}$, $\dot{\mu}_1 = \dot{\mu}$. We implement numerical simulations to verify our theory. The numerical results are shown in Figure 1.

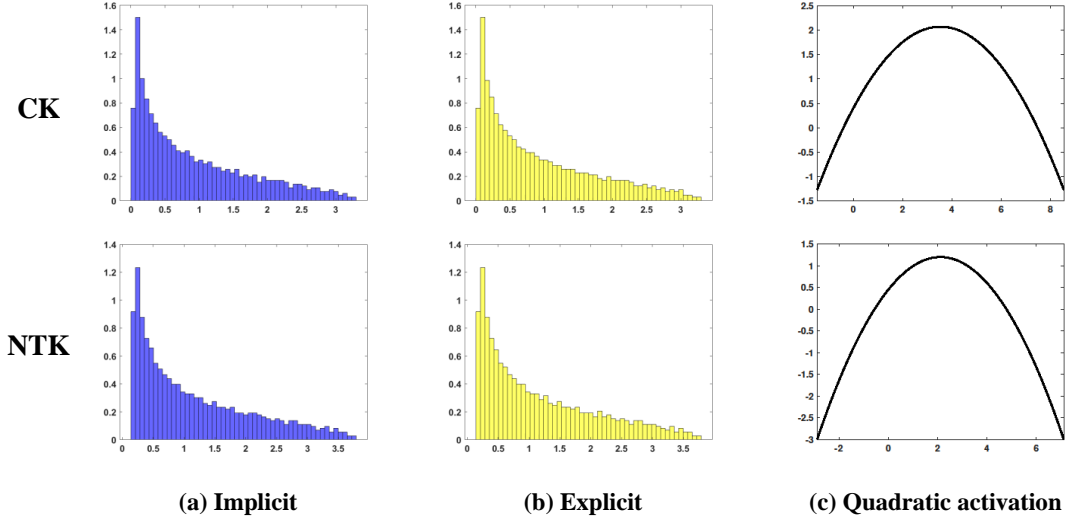


Figure 1: We independently generate $n = 1\,000$ data points from the $d = 1\,200$ -dimensional unit sphere. We use Gaussian initialization and σ_a^2 is set as 0.2. Upper: the CK results. Bottom: the NTK results. (a) spectral densities of implicit kernels, (b) spectral densities of explicit kernels, (c) quadratic activations.

Corollary 1 *We consider a quadratic polynomial activation $\sigma(t) = a_2 t^2 + a_1 t + a_0$. Let Assumptions 1 hold. As $n, d \rightarrow \infty$, the Implicit-CK matrix \mathbf{G}^* defined in Eq. (6) can be approximated consistently in operator norm, by the Explicit-CK matrix $\mathbf{\Sigma}$ defined in Eq. (5), i.e., $\|\mathbf{G}^* - \mathbf{\Sigma}\|_2 \rightarrow 0$, as long as*

$$a_2 = \pm \sqrt{\frac{\mu}{2}} \quad a_1 = \pm \sqrt{\beta}, \quad a_0 = \pm \sqrt{\alpha - \frac{\mu}{d}} - a_2,$$

and the Implicit-NTK matrix \mathbf{K}^ defined in Eq. (7) can be approximated consistently in operator norm, by the Explicit-NTK matrix $\mathbf{\Theta}$ defined in Eq. (5), i.e., $\|\mathbf{K}^* - \mathbf{\Theta}\|_2 \rightarrow 0$, as long as*

$$a_2 = \pm \sqrt{\frac{\dot{\mu}}{6}}, \quad a_1 = \pm \sqrt{\frac{\dot{\beta}}{2}}, \quad a_0 = \pm \sqrt{\dot{\alpha} - \frac{\dot{\mu}}{d}} - a_2.$$

4. Conclusion

In this paper, we study the CKs and NTKs of high-dimensional ReLU implicit NNs. We prove the asymptotic spectral equivalents for Implicit-CKs and Implicit-NTKs. Moreover, we establish the equivalence between implicit and explicit NNs by matching the coefficients of the asymptotic spectral equivalents. In particular, we show that a single-layer explicit NN with carefully designed activations has the same CK or NTK eigenspectra as a ReLU implicit NN. For future work, it would be interesting to extend our analysis to more general data distributions and activation functions.

Acknowledgements Z. Liao would like to acknowledge the National Natural Science Foundation of China (via fund NSFC-62206101) and the Fundamental Research Funds for the Central Universities of China (2021XXJS110) for providing partial support. R. C. Qiu and Z. Liao would like to acknowledge the National Natural Science Foundation of China (via fund NSFC-12141107), the Key Research and Development Program of Hubei (2021BAA037) and of Guangxi (GuiKe-AB21196034).

References

- [1] Hafiz Tiomoko Ali, Zhenyu Liao, and Romain Couillet. Random matrices in service of ml footprint: ternary random features with no performance loss. *ICLR*, 2022.
- [2] Shaojie Bai, J. Zico Kolter, and Vladlen Koltun. Deep equilibrium models. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [3] Shaojie Bai, Vladlen Koltun, and J Zico Kolter. Multiscale deep equilibrium models. *Advances in Neural Information Processing Systems*, 2020.
- [4] Noureddine El Karoui. The spectrum of kernel random matrices. *The Annals of Statistics*, 2010.
- [5] Zhou Fan and Zhichao Wang. Spectra of the conjugate kernel and neural tangent kernel for linear-width neural networks. *Advances in neural information processing systems*, 33:7710–7721, 2020.
- [6] Zhili Feng and J Zico Kolter. On the neural tangent kernel of equilibrium models. *arxiv*, 2020.
- [7] Lingyu Gu, Yongqi Du, Zhang Yuan, Di Xie, Shiliang Pu, Robert Qiu, and Zhenyu Liao. ” lossless” compression of deep neural networks: A high-dimensional neural tangent kernel approach. *Advances in Neural Information Processing Systems*, 35:3774–3787, 2022.
- [8] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.
- [9] Zenan Ling, Xingyu Xie, Qiu hao 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.
- [10] Lan V Truong. Global convergence rate of deep equilibrium models with general activations. *arXiv preprint arXiv:2302.05797*, 2023.
- [11] Xingyu Xie, Qiu hao 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*, 2022.