

GENERALIZATION OF SCALED DEEP RESNETS IN THE MEAN-FIELD REGIME

Yihang Chen

EPFL

yihang.chen@epfl.ch

Fanghui Liu

University of Warwick

fanghui.liu@warwick.ac.uk

Yiping Lu

New York University

yplu@nyu.edu

Grigorios G. Chrysos

University of Wisconsin-Madison

chrysos@wisc.edu

Volkan Cevher

EPFL

volkan.cevher@epfl.ch

ABSTRACT

Despite the widespread empirical success of ResNet, the generalization properties of deep ResNet are rarely explored beyond the lazy training regime. In this work, we investigate *scaled* ResNet in the limit of infinitely deep and wide neural networks, of which the gradient flow is described by a partial differential equation in the large-neural network limit, i.e., the *mean-field* regime. To derive the generalization bounds under this setting, our analysis necessitates a shift from the conventional time-invariant Gram matrix employed in the lazy training regime to a time-variant, distribution-dependent version. To this end, we provide a global lower bound on the minimum eigenvalue of the Gram matrix under the mean-field regime. Besides, for the traceability of the dynamic of Kullback-Leibler (KL) divergence, we establish the linear convergence of the empirical error and estimate the upper bound of the KL divergence over parameters distribution. Finally, we build the uniform convergence for generalization bound via Rademacher complexity. Our results offer new insights into the generalization ability of deep ResNet beyond the lazy training regime and contribute to advancing the understanding of the fundamental properties of deep neural networks.

1 INTRODUCTION

Deep neural networks (DNNs) have achieved great success empirically, a notable illustration of which is ResNet (He et al., 2016), a groundbreaking network architecture with skip connections. One typical way to theoretically understand ResNet (e.g., optimization, generalization), is based on the neural tangent kernel (NTK) tool (Jacot et al., 2018). Concretely, under proper assumptions, the training dynamics of ResNet can be described by a fixed kernel function (NTK). Hence, the global convergence and generalization guarantees can be given via NTK and the benefits of residual connection can be further demonstrated by spectral properties of NTK (Hayou et al., 2019; Huang et al., 2020; Hayou et al., 2021; Tirer et al., 2022). However, the NTK analysis requires the parameters of ResNet to not move much during training (which is called *lazy training* or kernel regime (Chizat et al., 2019; Woodworth et al., 2020; Barzilai et al., 2022)). Accordingly, the NTK analysis fails to describe the true non-linearity of ResNet. Beyond the NTK analysis thus received great attention in the deep learning theory community.

One typical approach is, *mean-field* based analysis, which allows for unrestricted movement of the parameters of DNNs during training, (Wei et al., 2019; Woodworth et al., 2020; Ghorbani et al., 2020; Yang & Hu, 2021; Akiyama & Suzuki, 2022; Ba et al., 2022; Mahankali et al., 2023; Chen et al., 2022; Mei et al., 2018; Rotskoff & Vanden-Eijnden, 2018; Nguyen, 2019; Sirignano & Spiliopoulos, 2020b). The training dynamics can be formulated as an optimization problem over the weight space of probability measures by studying suitable scaling limits. For deep ResNet, taking the limit of infinite depth is naturally connected to a continuous neural ordinary differential equation (ODE) (Sonoda & Murata, 2019; Weinan, 2017; Lu et al., 2018; Haber & Ruthotto, 2017; Chen et al., 2018), which makes the optimization guarantees of deep ResNets feasible under the mean-field regime (Lu et al.,

2020; Ma et al., 2020; Ding et al., 2021; 2022; Barboni et al., 2022). The obtained optimization results indicate that infinitely deep and wide ResNets can easily fit the training data with random labels, i.e., *global convergence*. While previous works have obtained promising optimization results for deep ResNets, there is a notable absence of generalization analysis, which is essential for theoretically understanding why deep ResNet can generalize well *beyond the lazy training regime*. Accordingly, this naturally raises the following question:

Can we build a generalization analysis of trained Deep ResNets in the mean-field setting?

We answer this question by providing a generalization analysis framework of *scaled* deep ResNet in the mean-field regime, where the word *scaled* denotes a scaling factor on deep ResNet. To this end, we consider an infinite width and infinite depth ResNet parameterized by two measures: ν over the feature encoder and τ over the output layer, respectively. By proving the condition number of the optimization dynamics of deep ResNets parameterized by τ and ν is lower bounded, we obtain the global linear convergence guarantees, aiming to derive the Kullback–Leibler (KL) divergence of such two measures between initialization and after training. Based on the KL divergence results, we, therefore, build the uniform convergence result for generalization via Rademacher complexity under the mean-field regime, obtaining the convergence rate at $\mathcal{O}(1/\sqrt{n})$ given n training data.

Our contributions are summarized as below:

- The paper provides the minimum eigenvalue estimation (lower bound) of the Gram matrix of the gradients for deep ResNet parameterized by the ResNet encoder’s parameters and MLP predictor’s parameters.
- The paper proves that the KL divergence of feature encoder τ and output layer ν can be bounded by a constant (depending only on network architecture parameters) during the training.
- This paper builds the connection between the Rademacher complexity result and KL divergence, and then derives the convergence rate $\mathcal{O}(1/\sqrt{n})$ for generalization.

Our theoretical analysis provides an in-depth understanding of the global convergence under minimal assumptions, sheds light on the KL divergence of network measures before and after training, and builds the generalization guarantees under the mean-field regime, matching classical results in the lazy training regime (Allen-Zhu et al., 2019; Du et al., 2019b). We expect that our analysis opens the door to generalization analysis for feature learning and look forward to which adaptive features can be learned from the data under the mean-field regime.

2 RELATED WORK

In this section, we briefly introduce the large width/depth ResNets in an ODE formulation, NTK analysis, and mean-field analysis for ResNets.

2.1 INFINITE-WIDTH, INFINITE-DEPTH RESNET, ODE

The limiting model of deep and wide ResNets can be categorized into three class, by either taking the width or depth to infinity: a) the infinite depth limit to the ODE/SDE model (Sonoda & Murata, 2019; Weinan, 2017; Lu et al., 2018; Chen et al., 2018; Haber & Ruthotto, 2017; Marion et al., 2023); b) infinite width limit, (Hayou et al., 2021; Hayou & Yang, 2023; Frei et al., 2019); c) infinite depth width ResNets, mean-field ODE framework under the infinite depth width limit (Li et al., 2021; Lu et al., 2020; Ding et al., 2021; 2022; Barboni et al., 2022)

In this work, we are particularly interested in mean-field ODE formulation. The deep ResNets modeling by mean-field ODE formulation stems from (Lu et al., 2020), in which every residual block is regarded as a particle and the target is changed to optimize over the empirical distribution of particles. Sander et al. (2022) discusses the rationale behind such equivalence between discrete dynamics and continuous ODE for ResNet under certain cases. The global convergence is built under a modified cost function (Ding et al., 2021), and further improved by removing the regularization term on the cost function (Ding et al., 2022). However, the analysis in (Ding et al., 2022) requires more technical assumptions about the limiting distribution. Barboni et al. (2022) show a local linear convergence by parameterizing the network with RKHS. However, the radius of the ball containing

parameters relies on the N -universality and is difficult to estimate. Our work requires minimal assumptions under a proper scaling of the network parameters and the design of architecture, and hence foresters optimization and generalization analyses. There is a concurrent works (Marion et al., 2023) that studies the implicit regularization of ResNets converging to ODEs, but the employed technique is different from ours and the generalization analysis in their work is missing.

2.2 NTK ANALYSIS FOR DEEP RESNET

Jacot et al. (2018) demonstrate that the training process of wide neural networks under gradient flow can be effectively described by the Neural Tangent Kernel (NTK) as the network’s width (denoted as M) tends towards infinity under the NTK scaling (Du et al., 2019b). During the training, the NTK remains unchanged and thus the theoretical analyses of neural networks can be transformed to those of kernel methods. In this case, the optimization and generalization properties of neural networks can be well controlled by the minimum eigenvalue of NTK (Cao & Gu, 2019; Nguyen et al., 2021). Regarding ResNets, the architecture we are interested in, the NTK analysis is also valid (Tirer et al., 2022; Huang et al., 2020; Belfer et al., 2021), as well as an algorithm-dependent bound (Frei et al., 2019) in the lazy training regime. Compared with kernelized analysis of the wide neural network, we do not rely on the convergence to a fixed kernel as the width approaches infinity to perform the convergence and generalization analysis. Instead, under the mean field regime, the so-called kernel falls into a time-varying, measure-dependent version. We also remark that, for a ResNet with an infinite depth but constant width, the global convergence is given by Cont et al. (2022) beyond the lazy training regime by studying the evolution of gradient norm. To our knowledge, this is the first work that analyzes the (varying) kernel eigenvalue of infinite-width/depth ResNet beyond the NTK regime in terms of optimization and generalization.

2.3 MEAN-FIELD ANALYSIS

Under suitable scaling limits (Mei et al., 2018; Rotskoff & Vanden-Eijnden, 2018; Sirignano & Spiliopoulos, 2020a), as the number of neurons goes to infinity, *i.e.* $M \rightarrow \infty$, neural networks work in the *mean-field limit*. In this setting, the training dynamics of neural networks can be formulated as an optimization problem over the distribution of neurons. A notable benefit of the mean-field approach is that, after deriving a formula for the gradient flow, conventional PDE methods can be utilized to characterize convergence behavior, which enables both *nonlinear feature learning* and global convergence (Araújo et al., 2019; Fang et al., 2019; Nguyen, 2019; Du et al., 2019a; Chatterji et al., 2021; Chizat & Bach, 2018; Mei et al., 2018; Wojtowysch, 2020; Lu et al., 2020; Sirignano & Spiliopoulos, 2021; 2020b; E et al., 2020; Jabir et al., 2021).

In the case of the two-layer neural network, Chizat & Bach (2018); Mei et al. (2018); Wojtowysch (2020); Chen et al. (2020); Barboni et al. (2022) justify the mean-field approach and demonstrate the convergence of the gradient flow process, achieving the zero loss. For the wide shallow neural network, Chen et al. (2022) proves the linear convergence of the training loss by virtue of the Gram matrix. In the multi-layer case, Lu et al. (2020); Ding et al. (2021; 2022) translate the training process of ResNet to a gradient-flow partial differential equation (PDE) and showed that with depth and width depending algebraically on the accuracy and confidence levels, first-order optimization methods can be guaranteed to find global minimizers that fit the training data.

In terms of the generalization of a trained neural network under the mean-field regime, current results are limited to two-layer neural networks. For example, Chen et al. (2020) provides a generalized NTK framework for two-layer neural networks, which exhibits a "kernel-like" behavior. Chizat & Bach (2020) demonstrate that the limits of the gradient flow of two-layer neural networks can be characterized as a max-margin classifier in a certain non-Hilbertian space. Our work, instead, focuses on deep ResNets in the mean-field regime and derives the generalization analysis framework.

3 FROM DISCRETE TO CONTINUOUS RESNET

In this section, we present the problem setting of our deep ResNets for binary classification under the infinite depth and width limit, which allows for parameter evolution of ResNets. Besides, several mild assumptions are introduced for our proof.

3.1 PROBLEM SETTING

For an integer L , we use the shorthand $[L] = \{1, 2, \dots, L\}$. Let $\mathcal{X} \subseteq \mathbb{R}^d$ be a compact metric space and $\mathcal{Y} \subseteq \mathbb{R}$. We assume that the training set $\mathcal{D}_n = \{(\mathbf{x}_i, y_i)\}_{i=1}^n$ is drawn from an unknown distribution μ on $\mathcal{X} \times \mathcal{Y}$, and μ_X is the marginal distribution of μ over \mathcal{X} . The goal of our supervised learning task is to find a hypothesis (i.e., a ResNet used in this work) $f : \mathcal{X} \rightarrow \mathcal{Y}$ such that $f(\mathbf{x}; \Theta)$ parameterized by Θ is a good approximation of the label $y \in \mathcal{Y}$ corresponding to a new sample $\mathbf{x} \in \mathcal{X}$. In this paper, we consider a binary classification task, denoted by minimizing the expected risk, let $\ell_{0-1}(f, y) = \mathbb{1}\{yf < 0\}$,

$$\min_{\Theta} L_{0-1}(\Theta) := \mathbb{E}_{(\mathbf{x}, y) \sim \mu} \ell_{0-1}(f(\mathbf{x}; \Theta), y).$$

Note that the 0 – 1 loss is non-convex and non-smooth, and thus difficult for optimization. One standard way in practice for training is using a *surrogate* loss for empirical risk minimization (ERM), normally convex and smooth (or at least continuous), e.g., the hinge loss, the cross-entropy loss. Interestingly, the squared loss, originally used for regression, can be also applied for classification with good statistical properties in terms of robustness and calibration error, as systematically discussed in Hui & Belkin (2020); Hu et al. (2022). Therefore, we employ the squared loss in ERM for training, let $\ell(f, y) = \frac{1}{2}(y - f)^2$,

$$\min_{\Theta} \widehat{L}(\Theta) := \frac{1}{n} \sum_{i=1}^n \ell(f(\mathbf{x}_i; \Theta), y_i) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} \ell(f(\mathbf{x}; \Theta), y(\mathbf{x})), \quad (1)$$

where \mathcal{D}_n is the empirical measure of μ_X over $\{\mathbf{x}_i\}_{i=1}^n$ and note that y is a function of \mathbf{x} .

We call the probability measure $\rho \in \mathcal{P}^2$ if ρ has the finite second moment, and $\rho \in \mathcal{C}(\mathcal{P}^2; [0, 1])$ if $\rho^s \in \mathcal{P}^2, \forall s \in [0, 1]$. For $\rho_1, \rho_2 \in \mathcal{P}^2$, the 2-Wasserstein distance is denoted by $\mathcal{W}_2(\rho_1, \rho_2)$; and for $\rho_1, \rho_2 \in \mathcal{C}(\mathcal{P}^2; [0, 1])$, we define $\mathcal{W}_2(\rho_1, \rho_2) := \sup_{s \in [0, 1]} \mathcal{W}_2(\rho_1^s, \rho_2^s)$.

3.2 RESNETS IN THE INFINITE DEPTH AND WIDTH LIMIT

The continuous formulation of ResNets is a recent approach that uses differential equations to model the behavior of the ResNet. This formulation has the advantage of enabling continuous analysis of the ODE, which can make the analysis of ResNets easier (Lu et al., 2020; Ding et al., 2021; 2022; Barboni et al., 2022). We firstly consider the following ResNet (He et al., 2016) of depth L can be formulated as $z_0(\mathbf{x}) = \mathbf{x} \in \mathbb{R}^d$, and

$$\begin{aligned} z_{l+1}(\mathbf{x}) &= z_l(\mathbf{x}) + \frac{\alpha}{ML} \sum_{m=1}^M \sigma(z_l(\mathbf{x}), \boldsymbol{\theta}_{l,m}) \in \mathbb{R}^d, \quad l \in [L-1], \\ f_{\Omega_K, \Theta_{L,M}}(\mathbf{x}) &= \frac{\beta}{K} \sum_{k=1}^K h(z_L, \boldsymbol{\omega}_k) \in \mathbb{R}, \end{aligned} \quad (2)$$

where $\mathbf{x} \in \mathbb{R}^d$ is the input data, $\alpha, \beta \in \mathbb{R}_+$ are the scaling factors. $\Theta_{L,M} = \{\boldsymbol{\theta}_{l,m} \in \mathbb{R}^{k_\nu}\}_{l=0, m=0}^{L-1, M}$ is the parameters of the ResNet encoder $\sigma : \mathbb{R}^d \rightarrow \mathbb{R}^d$ (activation functions are implicitly included into σ), and $\Omega_K = \{\boldsymbol{\omega}_k \in \mathbb{R}^{k_\tau}\}_{k=1}^K$ is the parameters of the predictor $h : \mathbb{R}^d \rightarrow \mathbb{R}$. We introduce a trainable MLP parametrized by $\boldsymbol{\omega}$ in the end, which is different from the fixed linear predictor in Lu et al. (2020); Ding et al. (2021; 2022). We make the assumptions on the choices of activation function σ and predictor h later in Assumption 3.3. Different scaling of α, β leads to different training schemes. Note that setting $\alpha = \sqrt{M}, \beta = \sqrt{K}$ corresponds to the standard scaling in the NTK regime (Du et al., 2019b), while setting $\alpha = \beta = 1$ corresponds to the classical mean field analysis (Mei et al., 2018; Rotskoff & Vanden-Eijnden, 2018; Lu et al., 2020; Ding et al., 2022). We will keep α, β as a hyperparameter in our theoretical analysis and determine the choice of α, β in future discussions. Besides, the scaling $1/L$ is necessary to derive the neural ODE limit (Marion et al., 2022), which has been supported by the empirical observations from Bachlechner et al. (2021); Marion et al. (2023).

We then introduce the infinitely deep and wide ResNet which is known as the mean-field limit of deep ResNet (Lu et al., 2020; Ma et al., 2020; Ding et al., 2022).

Infinite Depth To be specific, we re-parametrize the indices $l \in [L]$ in Eq. (2) with $s = \frac{l}{L} \in [0, 1]$. We view z in Eq. (2) as a function in s that satisfies a coupled ODE, with $1/L$ being the stepsize. Accordingly, we write $\theta_m(s) := \theta_m(l/L) = \theta_{l,m}$, and $\Theta_M(s) = \{\theta_m(s)\}_{m=1}^M$. The continuous limit of Eq. (2) by taking $L \rightarrow \infty$ is

$$\frac{dz(\mathbf{x}, s)}{ds} = \frac{\alpha}{M} \sum_{m=1}^M \sigma(z(\mathbf{x}, s), \theta_m(s)) = \alpha \int_{\mathbb{R}^{k\nu}} \sigma(z(\mathbf{x}, s), \boldsymbol{\theta}) d\nu_M(\boldsymbol{\theta}, s), \quad z(\mathbf{x}, 0) = \mathbf{x}, \quad (3)$$

where the discrete probability $\nu_M(\boldsymbol{\theta}, s)$ is defined as $\nu_M(\boldsymbol{\theta}, s) := \frac{1}{M} \sum_{i=1}^M \delta_{\theta_m(s)}(\boldsymbol{\theta})$. Accordingly, the empirical risk in Eq. (1) can be written as

$$\widehat{L}(\boldsymbol{\Omega}_K, \boldsymbol{\Theta}_M) := \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} \ell(f_{\boldsymbol{\Omega}_K, \boldsymbol{\Theta}_M}(\mathbf{x}), y(\mathbf{x})). \quad (4)$$

Infinite Width The mean-field limit is obtained by considering a ResNet of infinite width, i.e. $M \rightarrow \infty$. Denoting the limiting density of $\nu_M(\boldsymbol{\theta}, s)$ by $\nu(\boldsymbol{\theta}, s) \in \mathcal{C}(\mathcal{P}^2; [0, 1])$, Eq. (3) can be written as

$$\frac{dz(\mathbf{x}, s)}{ds} = \alpha \cdot \int_{\mathbb{R}^{k\nu}} \sigma(z(\mathbf{x}, s), \boldsymbol{\theta}) d\nu(\boldsymbol{\theta}, s), \quad s \in [0, 1], \quad z(\mathbf{x}, 0) = \mathbf{x}. \quad (5)$$

We denote the solution of Eq. (5) as $\mathbf{Z}_\nu(\mathbf{x}, s)$. Besides, we also take the infinite width limit in the final layer, i.e. $K \rightarrow \infty$, and then the limiting density of $\boldsymbol{\omega}$ is $\tau(\boldsymbol{\omega})$. The whole network can be written as

$$f_{\tau, \nu}(\mathbf{x}) := \beta \cdot \int_{\mathbb{R}^{k\tau}} h(\mathbf{Z}_\nu(\mathbf{x}, 1), \boldsymbol{\omega}) d\tau(\boldsymbol{\omega}), \quad (6)$$

and the empirical loss in Eq. (1) can be defined as:

$$\widehat{L}(\tau, \nu) := \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} \ell(f_{\tau, \nu}(\mathbf{x}), y(\mathbf{x})). \quad (7)$$

3.2.1 PARAMETER EVOLUTION

In the discrete ResNet (2), consider minimizing the empirical loss $\widehat{L}(\boldsymbol{\Omega}_K, \boldsymbol{\Theta}_{L,M})$ with an infinitesimally small learning rate, the updating process can be characterized by the particle gradient flow, see Definition 2.2 in Chizat & Bach (2018):

$$\frac{d\boldsymbol{\Omega}_K(t)}{dt} = -K \nabla_{\boldsymbol{\Omega}_K} \widehat{L}(\boldsymbol{\Omega}_K(t), \boldsymbol{\Theta}_{L,M}(t)), \quad (8)$$

$$\frac{d\boldsymbol{\Theta}_{L,M}(t)}{dt} = -LM \nabla_{\boldsymbol{\Theta}_{L,M}} \widehat{L}(\boldsymbol{\Omega}_K(t), \boldsymbol{\Theta}_{L,M}(t)), \quad (9)$$

where t is the rescaled pseudo-time, which amounts to assigning a $\frac{1}{K}$ or $\frac{1}{LM}$ mass to each particle, and is convenient to take the many-particle limit.

In the continuous ResNet (5), we use the gradient flow in the Wasserstein metric to characterize the evolution of τ, ν (Chizat & Bach, 2018). The evolution of the final layer distribution $\tau(\boldsymbol{\omega})$ can be characterized as

$$\frac{\partial \tau}{\partial t}(\boldsymbol{\omega}, t) = \nabla_{\boldsymbol{\omega}} \cdot \left(\tau(\boldsymbol{\omega}, t) \nabla_{\boldsymbol{\omega}} \frac{\delta \widehat{L}(\tau, \nu)}{\delta \tau}(\boldsymbol{\omega}, t) \right), \quad t \geq 0, \quad (10)$$

where

$$\frac{\delta \widehat{L}(\tau, \nu)}{\delta \tau}(\boldsymbol{\omega}) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} [\beta \cdot (f_{\tau, \nu}(\mathbf{x}) - y(\mathbf{x})) \cdot h(\mathbf{Z}_\nu(\mathbf{x}, 1), \boldsymbol{\omega})]. \quad (11)$$

In addition, the evolution of the ResNet layer $\nu(\boldsymbol{\theta}, s)$ can be characterized as

$$\frac{\partial \nu}{\partial t}(\boldsymbol{\theta}, s, t) = \nabla_{\boldsymbol{\theta}} \cdot \left(\nu(\boldsymbol{\theta}, s, t) \nabla_{\boldsymbol{\theta}} \frac{\delta \widehat{L}(\tau, \nu)}{\delta \nu}(\boldsymbol{\theta}, s, t) \right), \quad t \geq 0. \quad (12)$$

From the results in Lu et al. (2020); Ding et al. (2022; 2021), we can compute the functional derivative as follows:

$$\frac{\delta \widehat{\mathcal{L}}(\tau, \nu)}{\delta \nu}(\boldsymbol{\theta}, s) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n}[\beta \cdot (f_{\tau, \nu}(\mathbf{x}) - y(\mathbf{x})) \cdot \boldsymbol{\omega}^\top \frac{\partial \mathbf{Z}_\nu(\mathbf{x}, 1)}{\partial \mathbf{Z}_\nu(\mathbf{x}, s)} \frac{\delta \mathbf{Z}_\nu(\mathbf{x}, s)}{\delta \nu}(\boldsymbol{\theta}, s)] \quad (13)$$

$$= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n}[\beta \cdot (f_{\tau, \nu}(\mathbf{x}) - y(\mathbf{x})) \cdot \mathbf{p}_\nu^\top(\mathbf{x}, s) \cdot \alpha \cdot \boldsymbol{\sigma}(\mathbf{Z}_\nu(\mathbf{x}, s), \boldsymbol{\theta})], \quad (14)$$

where $\mathbf{p}_\nu \in \mathbb{R}^{k_\nu}$, parameterized by \mathbf{x}, s, ν , is the solution to the following adjoint ODE, with initial condition dependent on τ :

$$\frac{d\mathbf{p}_\nu^\top}{ds}(\mathbf{x}, s) = -\alpha \cdot \mathbf{p}_\nu^\top(\mathbf{x}, s) \int_{\mathbb{R}^{k_\nu}} \nabla_{\mathbf{z}} \boldsymbol{\sigma}(\mathbf{Z}_\nu(\mathbf{x}, s), \boldsymbol{\theta}) d\nu(\boldsymbol{\theta}, s), \quad (15)$$

$$\mathbf{p}_\nu^\top(\mathbf{x}, 1) = \int_{\mathbb{R}^{k_\tau}} \nabla_{\mathbf{z}} h(\mathbf{Z}_\nu(\mathbf{x}, 1), \boldsymbol{\omega}) d\tau(\boldsymbol{\omega}). \quad (16)$$

For the linear ODE (15), we can directly obtain the explicit formula, $\mathbf{p}_\nu^\top(\mathbf{x}, s) = \mathbf{p}_\nu^\top(\mathbf{x}, 1) \mathbf{q}_\nu(\mathbf{x}, s)$, where $\mathbf{q}_\nu(\mathbf{x}, s)$ is the exponentially scaling matrix defined in Eq. (17). The correctness of solution (17) can be verified by taking the gradient w.r.t. s at both sides.

$$\mathbf{q}_\nu(\mathbf{x}, s) = \exp\left(\alpha \int_s^1 \int_{\mathbb{R}^{k_\nu}} \nabla_{\mathbf{z}} \boldsymbol{\sigma}(\mathbf{Z}_\nu(\mathbf{x}, s'), \boldsymbol{\theta}) d\nu(\boldsymbol{\theta}, s')\right). \quad (17)$$

3.3 ASSUMPTIONS

In the following, we use the upper subscript for ResNet ODE layer $s \in [0, 1]$, and the lower subscript for training time $t \in [0, +\infty)$. For example, $\tau_t(\boldsymbol{\omega}) := \tau(\boldsymbol{\omega}, t)$, and $\nu_t^s(\boldsymbol{\theta}) := \nu(\boldsymbol{\theta}, s, t)$. First, we assume the boundedness and second moment of the dataset by Assumption 3.1.

Assumption 3.1 (Assumptions on data). *We assume that for $\mathbf{x}_i \neq \mathbf{x}_j \sim \mu_X$, the following holds with probability 1,*

$$\|\mathbf{x}_i\|_2 = 1, |y(\mathbf{x}_i)| \leq 1, \langle \mathbf{x}_i, \mathbf{x}_j \rangle \leq C_{\max} < 1, \forall i, j \in [n].$$

Remark: The assumption, i.e., $\mathbf{x}_i, \mathbf{x}_j$ being not parallel, is attainable and standard in the analysis of neural networks (Du et al., 2019b; Zhu et al., 2022).

Second, we adopt the standard Gaussian initialization for distribution τ and ν .

Assumption 3.2 (Assumption on initialization). *The initial distribution τ_0, ν_0 is standard Gaussian: $(\tau_0, \nu_0)(\boldsymbol{\omega}, \boldsymbol{\theta}, s) \propto \exp\left(-\frac{\|\boldsymbol{\omega}\|_2^2 + \|\boldsymbol{\theta}\|_2^2}{2}\right), \forall s \in [0, 1]$.*

Next, we adopt the following assumption on activation $\boldsymbol{\sigma}, h$ in terms of formulation and smoothness. The widely used activation functions, such as Sigmoid, Tanh, satisfy this assumption.

Assumption 3.3 (Assumptions on activation $\boldsymbol{\sigma}, h$). *Let $\boldsymbol{\theta} := (\mathbf{u}, \mathbf{w}, b) \in \mathbb{R}^{k_\nu}$, where $\mathbf{u}, \mathbf{w} \in \mathbb{R}^{k_\nu}, b \in \mathbb{R}$, i.e. $k_\nu = 2d + 1$; $\boldsymbol{\omega} := (a, \mathbf{w}, b) \in \mathbb{R}^{k_\tau}$, where $\mathbf{w} \in \mathbb{R}^{k_\nu}, a, b \in \mathbb{R}$, i.e. $k_\tau = d + 2$.*

For any $\mathbf{z} \in \mathbb{R}^{k_\nu}$, we assume

$$\boldsymbol{\sigma}(\mathbf{z}, \boldsymbol{\theta}) = \mathbf{u}\sigma_0(\mathbf{w}^\top \mathbf{z} + b), \quad h(\mathbf{z}, \boldsymbol{\omega}) = a\sigma_0(\mathbf{w}^\top \mathbf{z} + b), \quad \sigma_0: \mathbb{R} \rightarrow \mathbb{R}. \quad (18)$$

In addition, we have the following assumption on σ_0 . $|\sigma_0(x)| \leq C_1 \max(|x|, 1), |\sigma_0'(x)| \leq C_1, |\sigma_0''(x)| \leq C_1$, and let $\mu_i(\sigma_0)$ be the i -th Hermite coefficient of σ_0 .

Based on our description of the evolution of deep ResNets and standard assumptions, we are ready to present our main results on optimization and generalization in the following section.

4 MAIN RESULTS

In this section, we derive a quantitative estimation of the convergence rate of optimizing the ResNet. Our main results are three-fold: a) minimum eigenvalue estimation of the Gram matrix during the training dynamics which controls the training speed; b) a quantitative estimation of KL divergence between the weight destruction of trained ResNet with initialization; c) Rademacher complexity generalization guarantees for the trained ResNet.

4.1 GRAM MATRIX AND MINIMUM EIGENVALUE

The training dynamics is governed by the Gram matrix of the coordinate tangent vectors to the functional derivatives. In this section, we bound the minimum eigenvalue of the Gram matrix of the gradients through the whole training dynamics, which controls the convergence of gradient flow.

In the lazy training regime (Jacot et al., 2018), the Gram matrix converges pointwisely to the NTK as the width approaches infinity. Hence one only needs to bound the Gram matrix’s minimum eigenvalue at initialization, and the global convergence rate under gradient descent is bounded by the minimum eigenvalue of NTK. In our setting under the mean-field regime, we also need similar Gram matrix/matrices to aid our proof on the training dynamics, but we do not rely on their convergence to the NTK when the width approaches infinity. Instead, we consider the Gram matrix of the limiting mean-field model Eq. (6).

For the ResNet parameter distribution ν , we define one Gram matrix $\mathbf{G}_1(\tau, \nu) \in \mathbb{R}^{n \times n}$ by

$$\mathbf{G}_1(\tau, \nu) = \int_0^1 \mathbf{G}_1(\tau, \nu, s) ds, \quad \mathbf{G}_1(\tau, \nu, s) = \mathbb{E}_{\boldsymbol{\theta} \sim \nu(\cdot, s)} \mathbf{J}_1(\tau, \nu, \boldsymbol{\theta}, s) \mathbf{J}_1(\tau, \nu, \boldsymbol{\theta}, s)^\top, \quad (19)$$

where the row vector of \mathbf{J}_1 is defined as

$$(\mathbf{J}_1(\tau, \nu, \boldsymbol{\theta}, s))_{i,\cdot} = \mathbf{p}_\nu^\top(\mathbf{x}_i, s) \nabla_{\boldsymbol{\theta}} \boldsymbol{\sigma}(\mathbf{Z}_\nu(\mathbf{x}_i, s), \boldsymbol{\theta}), \quad 1 \leq i \leq n,$$

where the dependence on τ on the right side of the equality is from the initial condition $\mathbf{p}_\nu^\top(\mathbf{x}, 1)$. We also define the Gram matrix for the MLP parameter distribution τ , $\mathbf{G}_2(\tau, \nu) \in \mathbb{R}^{n \times n}$ by

$$\mathbf{G}_2(\tau, \nu) = \mathbb{E}_{\boldsymbol{\omega} \sim \tau(\cdot)} \mathbf{J}_2(\nu, \boldsymbol{\omega}) \mathbf{J}_2(\nu, \boldsymbol{\omega})^\top, \quad (20)$$

where the row vector of \mathbf{J}_2 is defined as

$$(\mathbf{J}_2(\nu, \boldsymbol{\omega}))_{i,\cdot} = \nabla_{\boldsymbol{\omega}} h(\mathbf{Z}_\nu(\mathbf{x}_i, 1), \boldsymbol{\omega}), \quad 1 \leq i \leq n.$$

We characterize the training dynamics of the neural networks by the following theorem (the proof deferred to Appendix C.1), which demonstrates the relationship between the gradient flow of the loss and those of functional derivatives.

Theorem 4.1. *The training dynamics of $\widehat{\mathcal{L}}(\tau_t, \nu_t)$ can be written as:*

$$\frac{d\widehat{\mathcal{L}}(\tau_t, \nu_t)}{dt} = - \int_0^1 \int_{\mathbb{R}^{k_\nu}} \left\| \nabla_{\boldsymbol{\theta}} \frac{\delta \widehat{\mathcal{L}}(\tau_t, \nu_t)}{\delta \nu_t}(\boldsymbol{\theta}, s) \right\|_2^2 d\nu_t(\boldsymbol{\theta}, s) - \int_{\mathbb{R}^{k_\nu}} \left\| \nabla_{\boldsymbol{\omega}} \frac{\delta \widehat{\mathcal{L}}(\tau_t, \nu_t)}{\delta \tau_t}(\boldsymbol{\omega}) \right\|_2^2 d\tau_t(\boldsymbol{\omega}).$$

From the definition of functional derivatives $\frac{\delta \widehat{\mathcal{L}}(\tau_t, \nu_t)}{\delta \nu_t}(\boldsymbol{\theta}, s)$ and $\frac{\delta \widehat{\mathcal{L}}(\tau_t, \nu_t)}{\delta \tau_t}(\boldsymbol{\omega})$, we immediately obtain Proposition 4.2, an extension of Theorem 4.1, which demonstrates that the training dynamics can be controlled by the corresponding Gram matrices.

Proposition 4.2. *Let $\mathbf{b}_t = (f_{\tau_t, \nu_t}(\mathbf{x}_1) - y(\mathbf{x}_1), \dots, f_{\tau_t, \nu_t}(\mathbf{x}_n) - y(\mathbf{x}_n))$, using the Gram matrix defined in Eq. (19) and (20), the training dynamics of $\widehat{\mathcal{L}}(\tau_t, \nu_t)$ can be written as:*

$$\frac{d\widehat{\mathcal{L}}(\tau_t, \nu_t)}{dt} = - \frac{\beta^2}{n^2} \mathbf{b}_t^\top (\alpha^2 \mathbf{G}_1(\tau_t, \nu_t) + \mathbf{G}_2(\tau_t, \nu_t)) \mathbf{b}_t.$$

Our analysis mainly relies on the minimum eigenvalue of the Gram matrix, which is commonly used in the analysis of overparameterized neural network (Arora et al., 2019; Chen et al., 2020). The minimum eigenvalue of the Gram matrix controls the convergence rate of the gradient descent.

We remark that the Gram matrix $\mathbf{G}_1(\tau_t, \nu_t)$ is always positive semi-definite for any $t \geq 0$, and $\mathbf{G}_1(\tau_0, \nu_0) = \mathbf{0}_{n \times n}$. Therefore, we only need to bound the minimum eigenvalue of $\mathbf{G}_2(\tau_t, \nu_t)$. First, we present such result under initialization, i.e., the lower bound of $\lambda_{\min}(\mathbf{G}_2(\tau_0, \nu_0))$ by the following lemma. The proof is deferred to Appendix C.2.

Lemma 4.3. *Under Assumption 3.1, 3.2, 3.3, there exist a constant $\Lambda := \Lambda(d)$, only depending on the dimension d , such that $\lambda_{\min}[\mathbf{G}(\tau_0, \nu_0)]$ is lower bounded by*

$$\lambda_0 := \lambda_{\min}(\mathbf{G}(\tau_0, \nu_0)) \geq \lambda_{\min}(\mathbf{G}_2(\tau_0, \nu_0)) \geq \Lambda(d).$$

Remark: Using the stability of the ODE model, we derive the KL divergence by virtue of the structure of the ResNet, build the lower bound of $\lambda_{\min}(\mathbf{G}_2)$, and prove the global convergence. In fact, our results, e.g., global convergence, and KL divergence can also depend on \mathbf{G}_1 by taking $\Lambda(t) := \alpha^2 \lambda_{\min}[\mathbf{G}_1(t)] + \lambda_{\min}(\mathbf{G}_2)$ in Lemma C.5. Due to $\lambda_{\min}[\mathbf{G}_1(t)] \geq 0$ for any t , we only use $\lambda_{\min}(\mathbf{G}_2) \geq \Lambda(d)$ in Lemma 4.3 for proof simplicity. Our model degenerates to a two-layer neural network if the residual part is removed (can be regarded as an identity mapping).

Second, for τ, ν different from initialization τ_0, ν_0 , we first prove in the finite time $t < t_{\max}$, we have the minimum eigenvalue is lower bounded $\lambda_{\min}(\mathbf{G}_2(\tau_t, \nu_t)) \geq \lambda_0/2$. In the next, we choose a proper scaling of α, β , such that $t_{\max} = \infty$, so that we obtain a global guarantee.

The proof is deferred to Appendix C.2.

Lemma 4.4. *There exists r_{\max} , such that, for $\nu \in \mathcal{C}(\mathcal{P}^2; [0, 1])$ and $\tau \in \mathcal{P}^2$ satisfying $\max\{\mathcal{W}_2(\nu, \nu_0), \mathcal{W}_2(\tau, \tau_0)\} \leq r_{\max}$, we have $\lambda_{\min}(\mathbf{G}_2(\tau, \nu)) \geq \frac{\lambda_0}{2}$.*

Remark: The radius is defined as $r_{\max} := \min\left\{\sqrt{d}, \frac{\Lambda(d)}{4nC_G(d, \alpha)}\right\}$, where $\Lambda(d)$ is defined in Lemma 4.3 and $C_G(d, \alpha)$ is some constant depending on d and α , used for the uniform estimation of $\mathbf{G}_2(\tau, \nu)$ around its initialization, refer to Lemma C.2. We detail this in Appendix C.3.

Definition 4.5. *Define*

$$t_{\max} := \sup\{t_0, \text{ s.t. } \forall t \in [0, t_0], \max\{\mathcal{W}_2(\nu_t, \nu_0), \mathcal{W}_2(\tau_t, \tau_0)\} \leq r_{\max}\},$$

where r_{\max} is defined in Lemma 4.4.

4.2 KL DIVERGENCE BETWEEN TRAINED NETWORK AND INITIALIZATION

Based on our previous results on the minimum eigenvalue of the Gram matrix, we are ready to prove the global convergence of the empirical loss over the weight distributions τ and ν of ResNets, and well control the KL divergence of them before and after training. The proofs in this subsection are deferred to Appendix C.4.

We first present the gradient flow of the KL divergence of the parameter distribution.

Lemma 4.6. *The dynamics of the KL divergence $\text{KL}(\tau_t \parallel \tau_0), \text{KL}(\nu_t \parallel \nu_0)$ through training can be characterize by*

$$\begin{aligned} \frac{d\text{KL}(\tau_t \parallel \tau_0)}{dt} &:= - \int_{\mathbb{R}^{k_\tau}} \left(\nabla_{\omega} \frac{\delta \text{KL}(\tau_t \parallel \tau_0)}{\delta \tau_t} \right) \cdot \left(\nabla_{\omega} \frac{\delta \widehat{\mathcal{L}}(\tau_t, \nu_t)}{\delta \tau_t} \right) d\tau_t(\omega), \\ \frac{d\text{KL}(\nu_t \parallel \nu_0)}{dt} &:= - \int_{\mathbb{R}^{k_\nu} \times [0, 1]} \left(\nabla_{\theta} \frac{\delta \text{KL}(\nu_t^s \parallel \nu_0^s)}{\delta \nu_t^s} \right) \cdot \left(\nabla_{\theta} \frac{\delta \widehat{\mathcal{L}}(\tau_t, \nu_t)}{\delta \nu_t} \right) d\nu_t(\theta, s). \end{aligned}$$

Since the evolution of τ_t, ν_t is continuous, we define the notation $t_{\max} > 0$ in the following, such that the minimum eigenvalue of $\mathbf{G}_2(\tau_t, \nu_t)$ can be lower bounded for $t < t_{\max}$. In our later proof, we will demonstrate that $t_{\max} = \infty$ can be achieved under proper α and β .

Combining Lemma 4.4 and Definition 4.5, we immediately obtain that $\lambda_{\min}(\mathbf{G}(\tau_t, \nu_t)) \geq \lambda_{\min}(\mathbf{G}_2(\tau_t, \nu_t)) \geq \lambda_0/2$ for $t < t_{\max}$.

By choosing certain α, β , we could prove $t_{\max} = \infty$, which leads to the bound $\text{KL}(\tau_t \parallel \tau_0)$, and $\text{KL}(\nu_t \parallel \nu_0)$ uniformly for all $t > 0$. (c.f. Theorem 4.7).

Theorem 4.7. *Assume the PDE (10) has solution $\tau_t \in \mathcal{P}^2$, and the PDE (12) has solution $\nu_t \in \mathcal{C}(\mathcal{P}^2; [0, 1])$. Under Assumption 3.1, 3.2, 3.3, for some constant C_{KL} dependent on d, α , taking $\bar{\beta} := \frac{\beta}{n} > \frac{4\sqrt{C_{\text{KL}}(d, \alpha)}}{\Lambda r_{\max}}$, the following results hold for all $t \in [0, \infty)$:*

$$\widehat{\mathcal{L}}(\tau_t, \nu_t) \leq e^{-\frac{\beta^2 \Lambda}{2n} t} \widehat{\mathcal{L}}(\tau_0, \nu_0), \quad \text{KL}(\tau_t \parallel \tau_0) \leq \frac{C_{\text{KL}}(d, \alpha)}{\Lambda^2 \bar{\beta}^2}, \quad \text{KL}(\nu_t \parallel \nu_0) \leq \frac{C_{\text{KL}}(d, \alpha)}{\Lambda^2 \bar{\beta}^2}.$$

We also derive a lower bound for the KL divergence. In Lemma C.10, we have that the average movement of the KL divergence is in the same order as the change in output layers.

4.3 RADEMACHER COMPLEXITY BOUND

After our previous estimates on the minimum eigenvalue of the Gram matrix as well as the KL divergence, we are ready to build the generalization bound for such trained mean-field ResNets. The proofs in this subsection are deferred to Appendix C.5.

Before we start the proof, we introduce some basic notations of Rademacher complexity. Let $\mathcal{D}_X = \{\mathbf{x}_i\}_{i=1}^n$ be the training dataset, and η_1, \dots, η_n be an i.i.d. family of Rademacher variables taking values ± 1 with equal probability. For any function set \mathcal{H} , the global Rademacher complexity is defined as $\mathcal{R}_n(\mathcal{H}) := \mathbb{E} [\sup_{h \in \mathcal{H}} \frac{1}{n} \sum_{i=1}^n \eta_i h(\mathbf{x}_i)]$.

Let $\mathcal{F} = \{f_{\tau, \nu}(\mathbf{x}) = \beta \cdot \int h(\mathbf{Z}_\nu(\mathbf{x}, 1), \boldsymbol{\omega}) d\tau(\boldsymbol{\omega})\}$ be the function class of infinite wide infinite depth ResNet defined in Section 3. We consider the following function class of infinite wide infinite depth ResNets whose KL divergence to the initial distribution is upper bounded by some $r > 0$: $\mathcal{F}_{\text{KL}}(r) = \{f_{\tau, \nu} \in \mathcal{F} : \text{KL}(\tau \parallel \tau_0) \leq r, \text{KL}(\nu \parallel \nu_0) \leq r\}$. The Rademacher complexity of $\mathcal{F}_{\text{KL}}(r)$ is given by the following lemma.

Lemma 4.8. *Under Assumption 3.3, if $r \leq r_0 = O(1/\sqrt{n})$, the Rademacher complexity of $\mathcal{F}_{\text{KL}}(r)$ can be bounded by $\mathcal{R}_n(\mathcal{F}_{\text{KL}}(r)) \lesssim \beta \sqrt{r/n}$, where \lesssim hides the constant dependence on d, α .*

Now we consider the generalization error of the 0-1 classification problem,

Theorem 4.9 (Generalization). *Assume $\tau_y \in \mathcal{C}(\mathcal{P}^2; [0, 1])$ and $\nu_y \in \mathcal{P}^2$ be the ground truth distributions, such that, $y(\mathbf{x}) = \mathbb{E}_{\boldsymbol{\omega} \sim \tau_y} h(\mathbf{Z}_{\nu_y}(\mathbf{x}, 1), \boldsymbol{\omega})$. Under the Assumption 3.1, 3.2 and 3.3, we set $\beta > \Omega(\sqrt{n})$. For any $\delta > 0$, with probability at least $1 - \delta$, the following bound holds:*

$$\mathbb{E}_{\mathbf{x} \sim \mu_X} \ell_{0-1}(f_{\tau_*, \nu_*}(\mathbf{x}), y(\mathbf{x})) \lesssim 1/\sqrt{n} + 6\sqrt{\log(2/\delta)/2n},$$

where \lesssim hides the constant dependence on d, α .

Remark: Our results of $O(1/\sqrt{n})$ matches the standard generalization error in the NTK regime (Du et al., 2019b). However, in contrast to setting $\alpha = \sqrt{M}, \beta = \sqrt{K}$ in Eq. (2) as the NTK regime, we directly analyze the ResNet in the limiting infinite width depth model in Eq. (6), and select proper choice of α, β independent of the width. We also validate our theoretical results by some numerical experiments in Appendix C.6.

5 CONCLUSION

In this paper, we build the generalization bound for trained deep results beyond the NTK regime under mild assumptions. Our results demonstrate that the KL divergence between the distribution of parameters after training and initialization of an infinitely width and deep ResNet can be controlled via lower bounding the eigenvalue of the Gram matrix during training. Under some stronger data assumptions, e.g., k -sparse parity problem (Suzuki et al., 2023), we may ensure that the limiting distribution of deep ResNet moves far away from its initialization in terms of KL divergence, which cannot be derived under the current setting. We leave it as the future work.

6 ACKNOWLEDGEMENT

This work was carried out in the EPFL LIONS group. This work was supported by Hasler Foundation Program: Hasler Responsible AI (project number 21043), the Army Research Office and was accomplished under Grant Number W911NF-24-1-0048, and Swiss National Science Foundation (SNSF) under grant number 200021_205011. Corresponding authors: Fanghui Liu and Yihang Chen.

REFERENCES

- Shunta Akiyama and Taiji Suzuki. Excess risk of two-layer relu neural networks in teacher-student settings and its superiority to kernel methods. *arXiv preprint arXiv:2205.14818*, 2022.
- Z. Allen-Zhu, Y. Li, and Z. Song. A convergence theory for deep learning via over-parameterization. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 242–252. PMLR, 09–15 Jun 2019.

- D. Araújo, R. Oliveira, and D. Yukimura. A mean-field limit for certain deep neural networks. *arXiv/1906.00193*, 2019.
- Sanjeev Arora, Simon Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *International Conference on Machine Learning*, pp. 322–332. PMLR, 2019.
- Jimmy Ba, Murat A Erdogdu, Taiji Suzuki, Zhichao Wang, Denny Wu, and Greg Yang. High-dimensional asymptotics of feature learning: How one gradient step improves the representation. *Advances in Neural Information Processing Systems*, 35:37932–37946, 2022.
- Thomas Bachlechner, Bodhisattwa Prasad Majumder, Henry Mao, Gary Cottrell, and Julian McAuley. Rezero is all you need: Fast convergence at large depth. In *Uncertainty in Artificial Intelligence*, pp. 1352–1361. PMLR, 2021.
- Raphaël Barboni, Gabriel Peyré, and François-Xavier Vialard. On global convergence of resnets: From finite to infinite width using linear parameterization. *Advances in Neural Information Processing Systems*, 35:16385–16397, 2022.
- Daniel Barzilai, Amnon Geifman, Meirav Galun, and Ronen Basri. A kernel perspective of skip connections in convolutional networks. *arXiv preprint arXiv:2211.14810*, 2022.
- Yuval Belfer, Amnon Geifman, Meirav Galun, and Ronen Basri. Spectral analysis of the neural tangent kernel for deep residual networks. *arXiv preprint arXiv:2104.03093*, 2021.
- Yuan Cao and Quanquan Gu. Generalization bounds of stochastic gradient descent for wide and deep neural networks. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- N. Chatterji, P. Long, and P. Bartlett. When does gradient descent with logistic loss interpolate using deep networks with smoothed relu activations? *arXiv/2102.04998*, 2021.
- Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- Zhengdao Chen, Eric Vanden-Eijnden, and Joan Bruna. On feature learning in neural networks with global convergence guarantees. *arXiv preprint arXiv:2204.10782*, 2022.
- Zixiang Chen, Yuan Cao, Quanquan Gu, and Tong Zhang. A generalized neural tangent kernel analysis for two-layer neural networks. *Advances in Neural Information Processing Systems*, 33: 13363–13373, 2020.
- L. Chizat and F. Bach. On the global convergence of gradient descent for over-parameterized models using optimal transport. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- Lenaic Chizat and Francis Bach. Implicit bias of gradient descent for wide two-layer neural networks trained with the logistic loss. In *Conference on Learning Theory*, pp. 1305–1338. PMLR, 2020.
- Lenaic Chizat, Edouard Oyallon, and Francis Bach. On lazy training in differentiable programming. *Advances in neural information processing systems*, 32, 2019.
- Rama Cont, Alain Rossier, and RenYuan Xu. Convergence and implicit regularization properties of gradient descent for deep residual networks. *arXiv preprint arXiv:2204.07261*, 2022.
- Zhiyan Ding, Shi Chen, Qin Li, and Stephen Wright. On the global convergence of gradient descent for multi-layer resnets in the mean-field regime. *arXiv preprint arXiv:2110.02926*, 2021.
- Zhiyan Ding, Shi Chen, Qin Li, and Stephen J Wright. Overparameterization of deep resnet: Zero loss and mean-field analysis. *J. Mach. Learn. Res.*, 23:48–1, 2022.
- Monroe D Donsker and SR Srinivasa Varadhan. Asymptotic evaluation of certain markov process expectations for large time, i. *Communications on Pure and Applied Mathematics*, 28(1):1–47, 1975.

- S. Du, J. Lee, H. Li, L. Wang, and X. Zhai. Gradient descent finds global minima of deep neural networks. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 1675–1685, 09–15 Jun 2019a.
- S. Du, X. Zhai, B. Póczos, and A. Singh. Gradient descent provably optimizes over-parameterized neural networks. In *International Conference on Learning Representations*, 2019b.
- W. E. C. Ma, and L. Wu. Machine learning from a continuous viewpoint, i. *Science China Mathematics*, 63(11):2233–2266, Sep 2020.
- C. Fang, Y. Gu, W. Zhang, and T. Zhang. Convex formulation of overparameterized deep neural networks. *arXiv/1911/07626*, 2019.
- S. Frei, Y. Cao, and Q. Gu. Algorithm-dependent generalization bounds for overparameterized deep residual networks. In *NeurIPS*, 2019.
- Behrooz Ghorbani, Song Mei, Theodor Misiakiewicz, and Andrea Montanari. When do neural networks outperform kernel methods? *Advances in Neural Information Processing Systems*, 33: 14820–14830, 2020.
- Eldad Haber and Lars Ruthotto. Stable architectures for deep neural networks. *Inverse problems*, 34(1):014004, 2017.
- Soufiane Hayou and Greg Yang. Width and depth limits commute in residual networks. In *International Conference on Machine Learning*, 2023.
- Soufiane Hayou, Arnaud Doucet, and Judith Rousseau. Exact convergence rates of the neural tangent kernel in the large depth limit. *arXiv preprint arXiv:1905.13654*, 2019.
- Soufiane Hayou, Eugenio Clerico, Bobby He, George Deligiannidis, Arnaud Doucet, and Judith Rousseau. Stable resnet. In *International Conference on Artificial Intelligence and Statistics*, pp. 1324–1332. PMLR, 2021.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pp. 630–645. Springer, 2016.
- Tianyang Hu, Jun Wang, Wenjia Wang, and Zhenguo Li. Understanding square loss in training overparameterized neural network classifiers. *Advances in Neural Information Processing Systems*, 35:16495–16508, 2022.
- Kaixuan Huang, Yuqing Wang, Molei Tao, and Tuo Zhao. Why do deep residual networks generalize better than deep feedforward networks?—a neural tangent kernel perspective. In *Advances in Neural Information Processing Systems*, volume 33, pp. 2698–2709, 2020.
- Like Hui and Mikhail Belkin. Evaluation of neural architectures trained with square loss vs cross-entropy in classification tasks. *arXiv preprint arXiv:2006.07322*, 2020.
- J. Jabir, D. Šiška, and Ł. Szpruch. Mean-field neural odes via relaxed optimal control. *arxiv/1912.05475*, 2021.
- A. Jacot, F. Gabriel, and C. Hongler. Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- Mufan Li, Mihai Nica, and Dan Roy. The future is log-Gaussian: ResNets and their infinite-depth-and-width limit at initialization. In *Advances in Neural Information Processing Systems*, pp. 7852–7864, 2021.
- Y. Lu, C. Ma, Y. Lu, J. Lu, and L. Ying. A mean field analysis of deep ResNet and beyond: Towards provably optimization via overparameterization from depth. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119, pp. 6426–6436, 13–18 Jul 2020.
- Yiping Lu, Aoxiao Zhong, Quanzheng Li, and Bin Dong. Beyond finite layer neural networks: Bridging deep architectures and numerical differential equations. In *International Conference on Machine Learning*, pp. 3276–3285. PMLR, 2018.

- Chao Ma, Lei Wu, et al. Machine learning from a continuous viewpoint, i. *Science China Mathematics*, 63(11):2233–2266, 2020.
- Arvind Mahankali, Jeff Z Haochen, Kefan Dong, Margalit Glasgow, and Tengyu Ma. Beyond ntk with vanilla gradient descent: A mean-field analysis of neural networks with polynomial width, samples, and time. *arXiv preprint arXiv:2306.16361*, 2023.
- Pierre Marion, Adeline Fermanian, Gérard Biau, and Jean-Philippe Vert. Scaling resnets in the large-depth regime. *arXiv preprint arXiv:2206.06929*, 2022.
- Pierre Marion, Yu-Han Wu, Michael E. Sander, and Gérard Biau. Implicit regularization of deep residual networks towards neural odes, 2023.
- S. Mei, A. Montanari, and P. M. 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.
- P. M. Nguyen. Mean field limit of the learning dynamics of multilayer neural networks. *arXiv/1902.02880*, 2019.
- Quynh Nguyen, Marco Mondelli, and Guido F Montufar. Tight bounds on the smallest eigenvalue of the neural tangent kernel for deep relu networks. In *International Conference on Machine Learning*, pp. 8119–8129. PMLR, 2021.
- Quynh N Nguyen and Marco Mondelli. Global convergence of deep networks with one wide layer followed by pyramidal topology. *Advances in Neural Information Processing Systems*, 33: 11961–11972, 2020.
- Felix Otto and Cédric Villani. Generalization of an inequality by talagrand and links with the logarithmic sobolev inequality. *Journal of Functional Analysis*, 173(2):361–400, 2000.
- Michael Poli, Stefano Massaroli, Atsushi Yamashita, Hajime Asama, Jinkyoo Park, and Stefano Ermon. Torchdyn: implicit models and neural numerical methods in pytorch. In *Neural Information Processing Systems, Workshop on Physical Reasoning and Inductive Biases for the Real World*, volume 2, 2021.
- Yury Polyanskiy and Yihong Wu. Wasserstein continuity of entropy and outer bounds for interference channels. *IEEE Transactions on Information Theory*, 62(7):3992–4002, 2016.
- Grant Rotskoff and Eric Vanden-Eijnden. Parameters as interacting particles: long time convergence and asymptotic error scaling of neural networks. *Advances in neural information processing systems*, 31, 2018.
- Michael Sander, Pierre Ablin, and Gabriel Peyré. Do residual neural networks discretize neural ordinary differential equations? In *Advances in Neural Information Processing Systems*, pp. 36520–36532, 2022.
- J. Sirignano and K. Spiliopoulos. Mean field analysis of neural networks: A central limit theorem. *Stochastic Processes and their Applications*, 130(3):1820–1852, 2020a.
- J. Sirignano and K. Spiliopoulos. Mean field analysis of neural networks: A law of large numbers. *SIAM Journal on Applied Mathematics*, 80(2):725–752, 2020b.
- J. Sirignano and K. Spiliopoulos. Mean field analysis of deep neural networks. *Mathematics of Operations Research*, 2021. doi: 10.1287/moor.2020.1118.
- Sho Sonoda and Noboru Murata. Transport analysis of infinitely deep neural network. *The Journal of Machine Learning Research*, 20(1):31–82, 2019.
- Taiji Suzuki, Denny Wu, Kazusato Oko, and Atsushi Nitanda. Feature learning via mean-field langevin dynamics: classifying sparse parities and beyond. In *Advances in Neural Information Processing Systems*, 2023.
- Tom Tirer, Joan Bruna, and Raja Giryes. Kernel-based smoothness analysis of residual networks. In *Mathematical and Scientific Machine Learning*, pp. 921–954. PMLR, 2022.

- Colin Wei, Jason D Lee, Qiang Liu, and Tengyu Ma. Regularization matters: Generalization and optimization of neural nets vs their induced kernel. *Advances in Neural Information Processing Systems*, 32, 2019.
- Ee Weinan. A proposal on machine learning via dynamical systems. *Communications in Mathematics and Statistics*, 1(5):1–11, 2017.
- S. Wojtowytsch. On the convergence of gradient descent training for two-layer relu-networks in the mean field regime. *arXiv/2005/13530*, 2020.
- Blake Woodworth, Suriya Gunasekar, Jason D Lee, Edward Moroshko, Pedro Savarese, Itay Golan, Daniel Soudry, and Nathan Srebro. Kernel and rich regimes in overparametrized models. In *Conference on Learning Theory*, pp. 3635–3673. PMLR, 2020.
- Greg Yang and Edward J Hu. Tensor programs iv: Feature learning in infinite-width neural networks. In *International Conference on Machine Learning*, pp. 11727–11737. PMLR, 2021.
- Zhenyu Zhu, Fanghui Liu, Grigorios Chrysos, and Volkan Cevher. Generalization properties of nas under activation and skip connection search. *Advances in Neural Information Processing Systems*, 35:23551–23565, 2022.

A OVERVIEW OF APPENDIX

We give a brief overview of the appendix here.

- **Appendix B.** In Appendix B.1, we prove some lemmas that will be useful. In Appendix B.2, we provide the estimation of the activation function σ . In Appendix B.3, we provide the prior estimation of Z_ν, p_ν .
- **Appendix C** In Appendix C.1, we prove the gradient flow of $\frac{d\hat{J}}{dt}$ and $\frac{d\text{KL}}{dt}$. In Appendix C.2, we bound the minimal eigenvalue at initialization. In Appendix C.3, we bound the perturbation of minimal eigenvalue. In Appendix C.4, we bound the KL divergence in finite time, and choose scaling parameters to prove the main results in Theorem 4.7. In Appendix C.5, we bound the KL divergence and provided the generalization bound.
- **Appendix C.6** We provide the experimental verification.

B USEFUL ESTIMATIONS

B.1 USEFUL LEMMAS

Lemma B.1 (2-Wasserstein continuity for functions of quadratic growth, Proposition 1 in Polyanskiy & Wu (2016)). *Let μ, ν be two probability measures on \mathbb{R}^d with finite second moments, and let $g : \mathbb{R}^d \rightarrow \mathbb{R}$ be a C_1 function obeying*

$$\|\nabla g(w)\|_2 \leq c_1 \|w\|_2 + c_2, \forall w \in \mathbb{R}^d,$$

for some constants $c_1 > 0$ and $c_2 \geq 0$. Then

$$|\mathbb{E}_{w \sim \mu} g(w) - \mathbb{E}_{w \sim \nu} g(w)| \leq (c_1 \sigma + c_2) \mathcal{W}_2(\mu, \nu),$$

where $\sigma^2 = \max\{\mathbb{E}_{w \sim \mu} \|w\|_2^2, \mathbb{E}_{w \sim \nu} \|w\|_2^2\}$

Lemma B.2 (Corollary 2.1 in Otto & Villani (2000)). *The probability measure $\nu_0(\boldsymbol{\theta}) \propto \exp(-\frac{\|\boldsymbol{\theta}\|_2^2}{2})$ satisfies following Talagrand inequality (in short $T(\frac{1}{2})$) for any $\nu \in \mathcal{P}^2(\mathbb{R}^{k\nu})$*

$$\mathcal{W}_2^2(\nu, \nu_0) \leq 4\text{KL}(\nu \| \nu_0).$$

Lemma B.3 (Donsker-Varadhan representation (Donsker & Varadhan, 1975)). *Let μ, λ be probability measures on a measurable space (X, Σ) . For any bounded, Σ -measurable functions $\Phi : X \rightarrow \mathbb{R}$:*

$$\int_X \Phi d\mu \leq \text{KL}(\mu \| \lambda) + \log \int_X \exp(\Phi) d\lambda.$$

B.2 ESTIMATION OF σ

Lemma B.4 (Boundedness of $\sigma(\mathbf{z}, \boldsymbol{\theta})$). *Under Assumption 3.3, for $\mathbf{x} \in \mathbb{R}^d, \boldsymbol{\theta} \in \mathbb{R}^k$, we have*

$$\|\sigma(\mathbf{z}, \boldsymbol{\theta})\|_2 \leq C_\sigma (\|\mathbf{z}\|_2 + 1)(\|\boldsymbol{\theta}\|_2^2 + 1), \quad (21)$$

$$\|\nabla_{\mathbf{z}} \sigma(\mathbf{z}, \boldsymbol{\theta})\|_F \leq C_\sigma (\|\boldsymbol{\theta}\|_2^2 + 1), \quad (22)$$

$$\|\nabla_{\boldsymbol{\theta}} \sigma(\mathbf{z}, \boldsymbol{\theta})\|_F \leq C_\sigma (\|\mathbf{z}\|_2 + 1)(\|\boldsymbol{\theta}\|_2 + 1), \quad (23)$$

$$\|\Delta_{\boldsymbol{\theta}} \sigma(\mathbf{z}, \boldsymbol{\theta})\|_2 \leq C_\sigma (\|\mathbf{z}\|_2^2 + 1)(\|\boldsymbol{\theta}\|_2 + 1), \quad (24)$$

$$\|\nabla_{\boldsymbol{\theta}} (\nabla_{\boldsymbol{\theta}} \sigma(\mathbf{z}, \boldsymbol{\theta}) \cdot \boldsymbol{\theta})\|_F \leq C_\sigma (\|\boldsymbol{\theta}\|_2 + 1)(\|\mathbf{z}\|_2 + 1), \quad (25)$$

$$\|\nabla_{\boldsymbol{\theta}} \Delta_{\boldsymbol{\theta}} \sigma(\mathbf{z}, \boldsymbol{\theta})\|_F \leq C_\sigma (\|\boldsymbol{\theta}\|_2 + 1)(\|\mathbf{z}\|_2^3 + 1), \quad (26)$$

where Δ is the Laplace operator. Let $\boldsymbol{\sigma} = (\sigma_i)_{i=1}^d$, $(\nabla_{\mathbf{z}} \boldsymbol{\sigma})_{ij} = \nabla_{z_j} \sigma_i$, $(\nabla_{\boldsymbol{\theta}} \boldsymbol{\sigma})_{ij} = \nabla_{\theta_j} \sigma_i$, $(\Delta_{\boldsymbol{\theta}} \boldsymbol{\sigma})_i = \Delta_{\boldsymbol{\theta}} \sigma_i$, $(\nabla_{\boldsymbol{\theta}} \boldsymbol{\sigma} \cdot \boldsymbol{\theta})_{ij} = (\nabla_{\boldsymbol{\theta}} \boldsymbol{\sigma})_{ij} \theta_j$.

Proof of Lemma B.4. We prove the relations directly in the following:

$$\|\sigma(\mathbf{z}, \boldsymbol{\theta})\|_2 = \|\mathbf{u} \sigma_0(\mathbf{w}^\top \mathbf{z} + b)\|_2 \leq C_1 \|\mathbf{u}\|_2 |\mathbf{w}^\top \mathbf{z} + b| \leq C_1 (\|\mathbf{z}\|_2 + 1)(\|\boldsymbol{\theta}\|_2^2 + 1), \quad (27)$$

$$\|\nabla_{\mathbf{z}} \boldsymbol{\sigma}(\mathbf{z}, \boldsymbol{\theta})\|_F \leq \|\mathbf{u}\|_2 |\mathbf{w} \sigma_0'(\mathbf{w}^\top \mathbf{z} + b)| \leq \|\mathbf{u}\|_2 \cdot C_1 \|\mathbf{w}\|_2 \leq C_1 (\|\boldsymbol{\theta}\|_2^2 + 1). \quad (28)$$

We can write $\nabla_{\theta}\sigma(\mathbf{z}, \theta) \in \mathbb{R}^{d \times k}$ by

$$(\nabla_{\theta}\sigma(\mathbf{z}, \theta))_{ij} = \begin{cases} \sigma_0(\mathbf{w}^{\top}\mathbf{z} + b) & j = i, \\ 0, & j \neq i, 1 \leq j \leq d, \\ u_i z_{j-d} \sigma'_0(\mathbf{w}^{\top}\mathbf{z} + b) & d+1 \leq j \leq 2d, \\ u_i \sigma'_0(\mathbf{w}^{\top}\mathbf{z} + b) & j = 2d+1. \end{cases} \quad (29)$$

Therefore,

$$\begin{aligned} \|\nabla_{\theta}\sigma(\mathbf{z}, \theta)\|_F^2 &= \sum_{i=1}^d [\sigma_0^2(\mathbf{w}^{\top}\mathbf{z} + b) + u_i^2 (\sigma'_0(\mathbf{w}^{\top}\mathbf{z} + b))^2 (1 + \|\mathbf{z}\|_2^2)] \\ &\leq dC_1^2 (\|\mathbf{w}\|_2 \|\mathbf{z}\|_2 + b)^2 + C_1^2 \|u\|_2^2 (1 + \|\mathbf{z}\|_2^2) \\ &\leq 2dC_1^2 (\|\mathbf{w}\|_2^2 \|\mathbf{z}\|_2^2 + b^2) + C_1^2 \|u\|_2^2 (1 + \|\mathbf{z}\|_2^2) \\ &\leq 2dC_1^2 (1 + \|\mathbf{z}\|_2^2) (\|\mathbf{w}\|_2^2 + \|u\|_2^2 + b^2 + 1) \\ &= 2dC_1^2 (1 + \|\mathbf{z}\|_2^2) (1 + \|\theta\|_2^2) \leq 2dC_1^2 (1 + \|\mathbf{z}\|_2^2) (1 + \|\theta\|_2^2). \end{aligned}$$

Therefore,

$$\|\nabla_{\theta}\sigma(\mathbf{z}, \theta)\|_F \leq \sqrt{2d}C_1 (1 + \|\mathbf{z}\|_2) (1 + \|\theta\|_2). \quad (30)$$

For $i \in [d]$

$$|\Delta_{\theta}\sigma(\mathbf{z}, \theta)_i| = |u_i (1 + \|\mathbf{z}\|_2^2) \sigma''_0(\mathbf{w}^{\top}\mathbf{z} + b)| \leq C_1 \cdot u_i (\|\mathbf{z}\|_2^2 + 1). \quad (31)$$

Therefore,

$$\|\Delta_{\theta}\sigma(\mathbf{z}, \theta)\|_2 \leq C_1 (\|\mathbf{z}\|_2^2 + 1) \cdot \sqrt{\sum_{i=1}^d u_i^2} \leq C_1 (\|\mathbf{z}\|_2^2 + 1) \cdot (\|\theta\|_2 + 1). \quad (32)$$

By Eq. (29), we obtain

$$\langle \nabla_{\theta}\sigma(\mathbf{z}, \theta)_{i,\cdot}, \theta \rangle = u_i \sigma_0(\mathbf{w}^{\top}\mathbf{z} + b) + u_i (\mathbf{w}^{\top}\mathbf{z} + b) \sigma'_0(\mathbf{w}^{\top}\mathbf{z} + b), \quad 1 \leq i \leq d.$$

Hence,

$$(\nabla_{\theta}\langle \nabla_{\theta}\sigma(\mathbf{z}, \theta)_{i,\cdot}, \theta \rangle)_j = \begin{cases} \sigma_0(\mathbf{w}^{\top}\mathbf{z} + b) + (\mathbf{w}^{\top}\mathbf{z} + b) \sigma'_0(\mathbf{w}^{\top}\mathbf{z} + b) & j = i, \\ 0 & j \neq i, 1 \leq j \leq d, \\ u_i z_{j-d} (\sigma'_0(y) + (y\sigma'_0(y))')|_{y=\mathbf{w}^{\top}\mathbf{z}+b} & d+1 \leq j \leq 2d, \\ u_i (\sigma'_0(y) + (y\sigma'_0(y))')|_{y=\mathbf{w}^{\top}\mathbf{z}+b} & j = 2d+1. \end{cases}$$

By Assumption 3.3, we have

$$\begin{aligned} \|\nabla_{\theta}\langle \nabla_{\theta}\sigma(\mathbf{z}, \theta)_{i,\cdot}, \theta \rangle\|_2^2 &= (\sigma_0(\mathbf{w}^{\top}\mathbf{z} + b) + (\mathbf{w}^{\top}\mathbf{z} + b) \sigma'_0(\mathbf{w}^{\top}\mathbf{z} + b))^2 \\ &\quad + (u_i (\sigma'_0(y) + (y\sigma'_0(y))')|_{y=\mathbf{w}^{\top}\mathbf{z}+b})^2 (1 + \|\mathbf{z}\|_2^2) \\ &\leq [2C_1 (\|\theta\|_2 + 1) (\|\mathbf{z}\|_2 + 1)]^2 + 4u_i^2 C_1^2 (1 + \|\mathbf{z}\|_2^2) \end{aligned}$$

Hence,

$$\begin{aligned} \|\nabla_{\theta}\langle \nabla_{\theta}\sigma(\mathbf{z}, \theta), \theta \rangle\|_F^2 &= \sum_{i=1}^d \|\nabla_{\theta}\langle \nabla_{\theta}\sigma(\mathbf{z}, \theta)_{i,\cdot}, \theta \rangle\|_2^2 \\ &\leq 4dC_1^2 (\|\theta\|_2 + 1)^2 (\|\mathbf{z}\|_2 + 1)^2 + 4\|u\|_2^2 C_1^2 (\|\mathbf{z}\|_2 + 1)^2 \\ &\leq (4d + 4)C_1^2 (\|\theta\|_2 + 1)^2 (\|\mathbf{z}\|_2 + 1)^2 \end{aligned} \quad (33)$$

For the last part, by Eq. (31),

$$\nabla_{\theta} \Delta_{\theta} \sigma(\mathbf{z}, \theta)_{ij} = \begin{cases} (1 + \|\mathbf{z}\|_2^2) \sigma_0''(\mathbf{w}^\top \mathbf{z} + b) & j = i \\ 0 & j \neq i, 1 \leq j \leq d \\ u_i z_{j-d} (1 + \|\mathbf{z}\|_2^2) \sigma_0'''(\mathbf{w}^\top \mathbf{z} + b) & d+1 \leq j \leq 2d \\ u_i (1 + \|\mathbf{z}\|_2^2) \sigma_0'''(\mathbf{w}^\top \mathbf{z} + b) & j = 2d+1 \end{cases}$$

Therefore,

$$\begin{aligned} \|\nabla_{\theta} \Delta_{\theta} \sigma(\mathbf{z}, \theta)\|_F &\leq C_1 (1 + \|\mathbf{z}\|_2^2) \sqrt{\sum_{i=1}^d 1 + u_i^2 (1 + \|\mathbf{z}\|_2^2)} \\ &\leq \sqrt{d} C_1 (\|\mathbf{z}\|_2^2 + 1)^{1.5} (\|\theta\|_2 + 1) \leq 3\sqrt{d} C_1 (\|\theta\|_2 + 1) (\|\mathbf{z}\|_2^3 + 1) \end{aligned} \quad (34)$$

The last inequality is from that, for $x > 0$

$$x^3 + 1 = x^3 + \frac{1}{2} + \frac{1}{2} \geq \frac{3}{2^{\frac{3}{2}}} x, \quad x^3 + 1 = 1 + \frac{x^3}{2} + \frac{x^3}{2} \geq \frac{3}{2^{\frac{3}{2}}} x^2,$$

then, we have

$$(1 + x^2)^{\frac{3}{2}} \leq (1 + x^2)(1 + x) = 1 + x + x^2 + x^3 \leq (1 + x^3)(1 + \frac{2^{\frac{3}{2}}}{3}) < 3(1 + x^3).$$

From Eq. (27), Eq. (28), Eq. (30), Eq. (32), Eq. (33), and Eq. (34), taking $C_{\sigma}^1 = 4\sqrt{d}C_1$, the proof is finished. We defer the definition of C_{σ} later. \square

Lemma B.5 (Stability of $\sigma(\mathbf{z}, \theta)$). *Under Assumption 3.3, for $\mathbf{x} \in \mathbb{R}^d$, $\theta \in \mathbb{R}^k$, we have*

$$\|\sigma(\mathbf{z}_1, \theta) - \sigma(\mathbf{z}_2, \theta)\|_2 \leq C_{\sigma} \cdot (\|\theta\|_2^2 + 1) \|\mathbf{z}_1 - \mathbf{z}_2\|_2 \quad (35)$$

$$\|\nabla_{\mathbf{z}} \sigma(\mathbf{z}_1, \theta) - \nabla_{\mathbf{z}} \sigma(\mathbf{z}_2, \theta)\|_F \leq C_{\sigma} \cdot (\|\theta\|_2^2 + 1) \|\mathbf{z}_1 - \mathbf{z}_2\|_2 \quad (36)$$

$$\|\nabla_{\mathbf{z}} \sigma(\mathbf{z}, \theta_1) - \nabla_{\mathbf{z}} \sigma(\mathbf{z}, \theta_2)\|_F \leq C_{\sigma} \cdot (\|\theta_1\|_2 + \|\theta_2\|_2 + 1) \|\theta_1 - \theta_2\|_2 \quad (37)$$

$$\|\nabla_{\theta} \sigma(\mathbf{z}, \theta_1) - \nabla_{\theta} \sigma(\mathbf{z}, \theta_2)\|_F \leq C_{\sigma} \cdot (\|\mathbf{z}\|_2 + 1) \|\theta_1 - \theta_2\|_2 \quad (38)$$

$$\|\nabla_{\theta} \sigma(\mathbf{z}_1, \theta) - \nabla_{\theta} \sigma(\mathbf{z}_2, \theta)\|_F \leq C_{\sigma} \cdot (\|\theta\|_2^2 + 1) \|\mathbf{z}_1 - \mathbf{z}_2\|_2 \quad (39)$$

Proof of Lemma B.5. By the mean-value theorem, we have there exists $\epsilon \in [0, 1]$

$$\begin{aligned} \|\sigma(\mathbf{z}_1, \theta) - \sigma(\mathbf{z}_2, \theta)\|_2 &\leq \|\nabla_{\mathbf{z}} \sigma(\mathbf{z}_1 + \epsilon(\mathbf{z}_2 - \mathbf{z}_1), \theta)\|_F \|\mathbf{z}_1 - \mathbf{z}_2\|_2 \\ &\leq C_{\sigma}^1 \cdot (\|\theta\|_2^2 + 1) \|\mathbf{z}_1 - \mathbf{z}_2\|_2 \end{aligned}$$

Denote by $\theta = (\mathbf{u}, \mathbf{w}, b)$, we have

$$\begin{aligned} \|\nabla_{\mathbf{z}} \sigma(\mathbf{z}_1, \theta) - \nabla_{\mathbf{z}} \sigma(\mathbf{z}_2, \theta)\|_F &\leq \|\mathbf{u} \mathbf{w}^\top (\sigma_0'(\mathbf{w}^\top \mathbf{z}_1 + b) - \sigma_0'(\mathbf{w}^\top \mathbf{z}_2 + b))\|_F \\ &\leq C_{\sigma}^1 \cdot (\|\theta\|_2^2 + 1) \|\mathbf{z}_1 - \mathbf{z}_2\|_2 \end{aligned}$$

and

$$\begin{aligned} \|\nabla_{\mathbf{z}} \sigma(\mathbf{z}, \theta_1) - \nabla_{\mathbf{z}} \sigma(\mathbf{z}, \theta_2)\|_F &\leq \|\mathbf{u}_1 \mathbf{w}_1^\top \sigma_0'(\mathbf{w}_1^\top \mathbf{z} + b_1) - \mathbf{u}_2 \mathbf{w}_2^\top \sigma_0'(\mathbf{w}_2^\top \mathbf{z} + b_2)\|_F \\ &\leq C_{\sigma}^1 \|\mathbf{u}_1 \mathbf{w}_1^\top - \mathbf{u}_2 \mathbf{w}_2^\top\|_F \leq C_{\sigma}^1 \|(\mathbf{u}_1 - \mathbf{u}_2)(\mathbf{w}_1 - \mathbf{w}_2)^\top + (\mathbf{u}_1 - \mathbf{u}_2) \mathbf{w}_2^\top + \mathbf{u}_1 (\mathbf{w}_1 - \mathbf{w}_2)^\top\|_F \\ &\leq 2C_{\sigma}^1 (\|\theta_1\|_2 + \|\theta_2\|_2 + 1) \|\theta_1 - \theta_2\|_2 \end{aligned}$$

In the next, we have

$$\begin{aligned} &(\nabla_{\theta} \sigma(\mathbf{z}, \theta_1) - \nabla_{\theta} \sigma(\mathbf{z}, \theta_2))_{ij} \\ &= \begin{cases} \sigma_0(\mathbf{w}_1^\top \mathbf{z} + b_1) - \sigma_0(\mathbf{w}_2^\top \mathbf{z} + b_2) & j = i, \\ 0, & j \neq i, 1 \leq j \leq d, \\ u_i^1 z_{j-d} \sigma_0'(\mathbf{w}_1^\top \mathbf{z} + b_1) - u_i^2 z_{j-d} \sigma_0'(\mathbf{w}_2^\top \mathbf{z} + b_2) & d+1 \leq j \leq 2d, \\ u_i^1 \sigma_0'(\mathbf{w}_1^\top \mathbf{z} + b_1) - u_i^2 \sigma_0'(\mathbf{w}_2^\top \mathbf{z} + b_2) & j = 2d+1. \end{cases} \end{aligned}$$

and then,

$$(\nabla_{\theta}\sigma(\mathbf{z}, \boldsymbol{\theta}_1))_{ij} - (\nabla_{\theta}\sigma(\mathbf{z}, \boldsymbol{\theta}_2))_{ij} = \begin{cases} C_{\sigma}^1 \cdot \|\mathbf{w}_1 - \mathbf{w}_2\|_2 \|\mathbf{z}\|_2 & j = i, \\ 0, & j \neq i, 1 \leq j \leq d, \\ C_{\sigma}^1 \cdot |u_i^1 - u_i^2| z_{j-d} & d+1 \leq j \leq 2d, \\ C_{\sigma}^1 \cdot |u_i^1 - u_i^2| & j = 2d+1. \end{cases}$$

Therefore,

$$\|\nabla_{\theta}\sigma(\mathbf{z}, \boldsymbol{\theta}_1) - \nabla_{\theta}\sigma(\mathbf{z}, \boldsymbol{\theta}_2)\|_F \leq \sqrt{2d}C_{\sigma}^1(\|\mathbf{z}\|_2 + 1) \cdot \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_2.$$

Similarly, we have

$$\begin{aligned} & (\nabla_{\theta}\sigma(\mathbf{z}_1, \boldsymbol{\theta}))_{ij} - (\nabla_{\theta}\sigma(\mathbf{z}_2, \boldsymbol{\theta}))_{ij} \\ &= \begin{cases} \sigma_0(\mathbf{w}^{\top}\mathbf{z}_1 + b) - \sigma_0(\mathbf{w}^{\top}\mathbf{z}_2 + b) & j = i, \\ 0, & j \neq i, 1 \leq j \leq d, \\ u_i z_{j-d}^1 \sigma_0'(\mathbf{w}^{\top}\mathbf{z}_1 + b) - u_i z_{j-d}^2 \sigma_0'(\mathbf{w}^{\top}\mathbf{z}_2 + b) & d+1 \leq j \leq 2d, \\ u_i \sigma_0'(\mathbf{w}^{\top}\mathbf{z}_1 + b) - u_i \sigma_0'(\mathbf{w}^{\top}\mathbf{z}_2 + b) & j = 2d+1, \end{cases} \end{aligned}$$

and then

$$(\nabla_{\theta}\sigma(\mathbf{z}, \boldsymbol{\theta}_1))_{ij} - (\nabla_{\theta}\sigma(\mathbf{z}, \boldsymbol{\theta}_2))_{ij} = \begin{cases} C_{\sigma}^1 \cdot \|\mathbf{w}\|_2 \|\mathbf{z}_1 - \mathbf{z}_2\|_2 & j = i, \\ 0, & j \neq i, 1 \leq j \leq d, \\ C_{\sigma}^1 \cdot u_i |z_{j-d}^1 - z_{j-d}^2| & d+1 \leq j \leq 2d, \\ C_{\sigma}^1 \cdot u_i \|\mathbf{z}_1 - \mathbf{z}_2\|_2 \|\mathbf{w}\|_2 & j = 2d+1. \end{cases}$$

Therefore,

$$\|\nabla_{\theta}\sigma(\mathbf{z}_1, \boldsymbol{\theta}) - \nabla_{\theta}\sigma(\mathbf{z}_2, \boldsymbol{\theta})\|_F \leq \sqrt{d}C_{\sigma}^1(\|\boldsymbol{\theta}\|_2^2 + 1)\|\mathbf{z}_1 - \mathbf{z}_2\|_2$$

taking $C_{\sigma}^2 = \sqrt{2d}C_{\sigma}^1$, the proof is finished. \square

Combined with the estimation in the proofs of Lemma B.4 and Lemma B.5, we let

$$C_{\sigma} := 6d \cdot C_1 \quad (40)$$

B.3 PRIOR ESTIMATION OF ODE

The Lemma B.6 and Lemma B.7 establishes the boundedness and stability of \mathbf{Z}_{ν} and \mathbf{p}_{ν} with respect to ν .

Lemma B.6 (Boundedness and Stability of \mathbf{Z}_{ν}). *Suppose that Assumption 3.3 holds and that \mathbf{x} is in the support of \mathcal{X} . Suppose that $\nu_1, \nu_2 \in \mathcal{C}(\mathcal{P}^2; [0, 1])$ and $\mathbf{Z}_{\nu_1}, \mathbf{Z}_{\nu_2}$ are the corresponding unique solutions in Eq. (5). Then the following two bounds are satisfied for all $s \in [0, 1]$:*

$$\|\mathbf{Z}_{\nu_1}(\mathbf{x}, s)\|_2 \leq C_{\mathbf{Z}}(\|\nu_1\|_{\infty}^2; \alpha),$$

and

$$\|\mathbf{Z}_{\nu_1}(\mathbf{x}, s) - \mathbf{Z}_{\nu_2}(\mathbf{x}, s)\|_2 \leq C_{\mathbf{Z}}(\|\nu_1\|_{\infty}^2, \|\nu_2\|_{\infty}^2; \alpha) \cdot \mathcal{W}_2(\nu_1, \nu_2),$$

where $C_{\mathbf{Z}}(\|\nu_1\|_{\infty}^2, \|\nu_2\|_{\infty}^2; \alpha)$ is a constant depending only on $\|\nu_1\|_{\infty}^2, \|\nu_2\|_{\infty}^2$ and α , and for $\nu \in \mathcal{C}(\mathcal{P}^2; [0, 1])$, we denote by $\|\nu\|_{\infty}^2 := \sup_{s \in [0, 1]} \mathbb{E}_{\boldsymbol{\theta} \sim \nu(\cdot, s)} \|\boldsymbol{\theta}\|_2^2 < \infty$.

Proof of Lemma B.6. We firstly demonstrate that Eq. (5) has a unique \mathcal{C}_1 solution and then prove the boundedness of \mathbf{Z}_{ν} under different probability measures.

By Lemma B.5, we have

$$\begin{aligned} \left\| \int_{\mathbb{R}^k} (\sigma(\mathbf{z}_1, \boldsymbol{\theta}) - \sigma(\mathbf{z}_2, \boldsymbol{\theta})) d\nu_1(\boldsymbol{\theta}, s) \right\|_2 &\leq C_{\sigma} \|\mathbf{z}_1 - \mathbf{z}_2\|_2 \int_{\mathbb{R}^k} (\|\boldsymbol{\theta}\|_2^2 + 1) d\nu_1(\boldsymbol{\theta}, s) \\ &\leq C_{\sigma} \|\mathbf{z}_1 - \mathbf{z}_2\|_2 (\|\nu_1\|_{\infty}^2 + 1), \end{aligned} \quad (41)$$

which implies that $\int_{\mathbb{R}^k} \sigma(\mathbf{z}_1, \boldsymbol{\theta}) d\nu_1(\boldsymbol{\theta}, s)$ is locally Lipschitz. Combining this with the a-priori estimate, the ODE theory implies that Eq. (5) has a unique C_1 solution.

In the next, we aim to prove the boundedness of \mathbf{Z}_{ν} . For any $s \in [0, 1]$, by Eq. (21) in Lemma B.4, we have

$$\left\| \int_{\mathbb{R}^k} \sigma(\mathbf{z}, \boldsymbol{\theta}) d\nu_1(\boldsymbol{\theta}, s) \right\|_2 \leq \int_{\mathbb{R}^k} \|\sigma(\mathbf{z}, \boldsymbol{\theta})\|_2 d\nu_1(\boldsymbol{\theta}, s) \leq C_{\sigma}(\|\mathbf{z}\|_2 + 1) \int_{\mathbb{R}^k} (\|\boldsymbol{\theta}\|_2^2 + 1) d\nu_1(\boldsymbol{\theta}, s).$$

To prove the boundedness of \mathbf{Z}_{ν_1} , using Eq. (5) and Lemma B.4, we have

$$\begin{aligned} \frac{d\|\mathbf{Z}_{\nu_1}(\mathbf{x}, s)\|_2^2}{ds} &= 2\mathbf{Z}_{\nu_1}^{\top}(\mathbf{x}, s) \frac{d\mathbf{Z}_{\nu_1}(\mathbf{x}, s)}{ds} \\ &\leq 2\alpha C_{\sigma}(\|\mathbf{Z}_{\nu_1}(\mathbf{x}, s)\|_2^2 + \|\mathbf{Z}_{\nu_1}(\mathbf{x}, s)\|_2) \int_{\mathbb{R}^k} (\|\boldsymbol{\theta}\|_2^2 + 1) d\nu_1(\boldsymbol{\theta}, s) \\ &\leq 4\alpha C_{\sigma}(\|\mathbf{Z}_{\nu_1}(\mathbf{x}, s)\|_2^2 + 1) \int_{\mathbb{R}^k} (\|\boldsymbol{\theta}\|_2^2 + 1) d\nu_1(\boldsymbol{\theta}, s). \end{aligned}$$

By Grönwall's inequality, and $\mathbf{Z}_{\nu_1}(\mathbf{x}, 0) = \mathbf{x}$, we have

$$\begin{aligned} \|\mathbf{Z}_{\nu_1}(\mathbf{x}, s)\|_2 &\leq \exp\left(2\alpha C_{\sigma} \left(\int_0^1 \int_{\mathbb{R}^k} \|\boldsymbol{\theta}\|_2^2 d\nu_1(\boldsymbol{\theta}, s) + 1\right)\right) (\|\mathbf{x}\|_2 + 1) \\ &\leq \exp(2\alpha C_{\sigma}(\|\nu_1\|_{\infty}^2 + 1))(\|\mathbf{x}\|_2 + 1). \end{aligned}$$

By Assumption 3.1, $\|\mathbf{x}\|_2 \leq 1$, we thus have a priori estimate of \mathbf{Z}_{ν_1} . Let $C_{\mathbf{Z}}^1(\|\nu_1\|_{\infty}^2; \alpha) := 2\exp(2\alpha C_{\sigma}(\|\nu_1\|_{\infty}^2 + 1))$, we have $\|\mathbf{Z}_{\nu_1}(\mathbf{x}, s)\|_2 \leq C_{\mathbf{Z}}^1(\|\nu_1\|_{\infty}^2; \alpha)$.

In the next, to estimate the difference under different measures ν_1 and ν_2 , define

$$\boldsymbol{\delta}(\mathbf{x}, s) = \mathbf{Z}_{\nu_1}(\mathbf{x}, s) - \mathbf{Z}_{\nu_2}(\mathbf{x}, s),$$

and we can easily obtain

$$\begin{aligned} \frac{d\|\boldsymbol{\delta}(\mathbf{x}, s)\|_2^2}{ds} &= 2\alpha \left\langle \boldsymbol{\delta}(\mathbf{x}, s), \int_{\mathbb{R}^k} \sigma(\mathbf{Z}_{\nu_1}(\mathbf{x}, s), \boldsymbol{\theta}) d\nu_1(\boldsymbol{\theta}, s) - \int_{\mathbb{R}^k} \sigma(\mathbf{Z}_{\nu_2}(\mathbf{x}, s), \boldsymbol{\theta}) d\nu_2(\boldsymbol{\theta}, s) \right\rangle \\ &:= 2\alpha \langle \boldsymbol{\delta}(\mathbf{x}, s), (\text{A}) + (\text{B}) \rangle, \end{aligned} \quad (42)$$

where, by Eq. (41), we have

$$\|(\text{A})\|_2 := \left\| \int_{\mathbb{R}^k} (\sigma(\mathbf{Z}_{\nu_1}(\mathbf{x}, s), \boldsymbol{\theta}) - \sigma(\mathbf{Z}_{\nu_2}(\mathbf{x}, s), \boldsymbol{\theta})) d\nu_1(\boldsymbol{\theta}, s) \right\|_2 \leq C_{\sigma} \|\boldsymbol{\delta}(\mathbf{x}, s)\|_2 (\|\nu_1\|_{\infty}^2 + 1),$$

and

$$(\text{B}) := \int_{\mathbb{R}^k} \sigma(\mathbf{Z}_{\nu_2}(\mathbf{x}, s), \boldsymbol{\theta}) d\nu_1(\boldsymbol{\theta}, s) - \int_{\mathbb{R}^k} \sigma(\mathbf{Z}_{\nu_2}(\mathbf{x}, s), \boldsymbol{\theta}) d\nu_2(\boldsymbol{\theta}, s).$$

By Lemma B.1 and $\|\nabla_{\boldsymbol{\theta}} \sigma(\mathbf{Z}_{\nu_2}(\mathbf{x}, s), \boldsymbol{\theta})\|_F \leq C_{\sigma} \cdot (\|\mathbf{Z}_{\nu_2}(\mathbf{x}, s), \boldsymbol{\theta}\|_2 + 1)(\|\boldsymbol{\theta}\|_2 + 1)$, we can bound (B) by

$$\begin{aligned} \|(\text{B})\|_2 &\leq C_{\sigma} \cdot (\|\mathbf{Z}_{\nu_2}(\mathbf{x}, s), \boldsymbol{\theta}\|_2 + 1) \cdot (\|\boldsymbol{\theta}\|_2 + 1) \cdot \mathcal{W}_2(\nu_1^s, \nu_2^s) \\ &\leq C_{\sigma} \cdot (\|\mathbf{Z}_{\nu_2}(\mathbf{x}, s), \boldsymbol{\theta}\|_2 + 1) \cdot (\sqrt{\max\{\|\nu_1^s\|_2^2, \|\nu_2^s\|_2^2\}} + 1) \cdot \mathcal{W}_2(\nu_1^s, \nu_2^s) \\ &\leq C_{\sigma} \cdot (C_{\mathbf{Z}}^1(\|\nu_2\|_{\infty}^2; \alpha) + 1) \cdot (\sqrt{\|\nu_1\|_{\infty}^2 + \|\nu_2\|_{\infty}^2} + 1) \mathcal{W}_2(\nu_1, \nu_2). \end{aligned}$$

Plugging the estimate of (A) and (B) into Eq. (42), we have

$$\begin{aligned} \frac{d\|\boldsymbol{\delta}(\mathbf{x}, s)\|_2^2}{ds} &\leq 2\alpha C_{\sigma} (\|\boldsymbol{\delta}(\mathbf{x}, s)\|_2^2 (\|\nu_1\|_{\infty}^2 + 1) \\ &\quad + \|\boldsymbol{\delta}(\mathbf{x}, s)\|_2 (C_{\mathbf{Z}}^1(\|\nu_1\|_{\infty}^2; \alpha) + 1) (\sqrt{\|\nu_1\|_{\infty}^2 + \|\nu_2\|_{\infty}^2} + 1) \cdot \mathcal{W}_2(\nu_1, \nu_2)) \\ &\leq 2\alpha C_{\sigma} (\|\boldsymbol{\delta}(\mathbf{x}, s)\|_2^2 + \mathcal{W}_2^2(\nu_1, \nu_2)) (\sqrt{\|\nu_1\|_{\infty}^2 + \|\nu_2\|_{\infty}^2} + 1)^2 (C_{\mathbf{Z}}^1(\|\nu_1\|_{\infty}^2; \alpha) + 1)^2. \end{aligned}$$

Since $\boldsymbol{\delta}(\mathbf{x}, 0) = 0$, by Grönwall's inequality, we have, $\forall s \in [0, 1]$,

$$\|\boldsymbol{\delta}(\mathbf{x}, s)\|_2 \leq (\exp(\alpha C_{\sigma}) - 1) \cdot (\sqrt{\|\nu_1\|_{\infty}^2 + \|\nu_2\|_{\infty}^2} + 1) (C_{\mathbf{Z}}^1(\|\nu_1\|_{\infty}^2; \alpha) + 1) \cdot \mathcal{W}_2(\nu_1, \nu_2),$$

which concludes the proof. \square

Lemma B.7 (Boundedness and Stability of \mathbf{p}_ν). *Suppose that Assumption 3.3 holds and that \mathbf{x} is in the support of \mathcal{X} . Suppose that $\nu_1, \nu_2 \in \mathcal{C}(\mathcal{P}^2; [0, 1])$ and $\mathbf{p}_{\nu_1}(\mathbf{x}, s)$, $\mathbf{p}_{\nu_2}(\mathbf{x}, s)$ are defined in Eq. (15). Then the following three bounds are satisfied for all $s \in [0, 1]$:*

$$\|\mathbf{p}_{\nu_1}(\mathbf{x}, s)\|_2 \leq C_{\mathbf{p}}(\|\nu_1\|_\infty^2, \|\tau\|_2^2; \alpha), \quad (43)$$

and

$$\|\mathbf{p}_{\nu_1}(\mathbf{x}, s) - \mathbf{p}_{\nu_2}(\mathbf{x}, s)\|_2 \leq C_{\mathbf{p}}(\|\nu_1\|_\infty^2, \|\nu_2\|_\infty^2, \|\tau\|_2^2; \alpha) \cdot \mathcal{W}_2(\nu_1, \nu_2), \quad (44)$$

where $C_{\mathbf{p}}(\|\nu_1\|_\infty^2, \|\nu_2\|_\infty^2, \|\tau\|_2^2; \alpha)$ is a constant depending only on $\|\nu_1\|_\infty^2, \|\nu_2\|_\infty^2, \|\tau\|_2^2$ and α , and for $\tau \in \mathcal{P}^2$, we denote by $\|\tau\|_2^2 := \mathbb{E}_{\omega \sim \tau(\cdot)} \|\tau\|_2^2 < \infty$.

Proof of Lemma B.7. For any $s \in [0, 1]$, by Eq. (15) and the estimation of $\nabla_{\mathbf{z}} \sigma$ in Eq. (22), we have

$$\begin{aligned} \frac{d\|\mathbf{p}_{\nu_1}(\mathbf{x}, s)\|_2^2}{ds} &= 2 \frac{d\mathbf{p}_{\nu_1}^\top(\mathbf{x}, s)}{ds} \mathbf{p}_{\nu_1}(\mathbf{x}, s) \\ &\leq 2\alpha \|\mathbf{p}_{\nu_1}^\top(\mathbf{x}, s)\|_2^2 \cdot \left\| \int_{\mathbb{R}^k} \nabla_{\mathbf{z}} \sigma(\mathbf{Z}_{\nu_1}(\mathbf{x}, s), \boldsymbol{\theta}) d\nu_1(\boldsymbol{\theta}, s) \right\|_F \\ &\leq 2\alpha C_\sigma \|\mathbf{p}_{\nu_1}^\top(\mathbf{x}, s)\|_2^2 \int_{\mathbb{R}^k} (\|\boldsymbol{\theta}\|_2^2 + 1) d\nu_1(\boldsymbol{\theta}, s). \end{aligned}$$

It follows from the estimation of $\nabla_{\mathbf{z}} \sigma$ in Eq. (22),

$$\|\mathbf{p}_{\nu_1}^\top(\mathbf{x}, 1)\|_2 = \left\| \int_{\mathbb{R}^{k\tau}} \nabla_{\mathbf{z}} h(\mathbf{Z}_{\nu_1}(\mathbf{x}, 1), \boldsymbol{\omega}) d\tau(\boldsymbol{\omega}) \right\|_2 \leq C_\sigma \cdot (\|\tau\|_2^2 + 1).$$

Therefore, by the Grönwall's inequality

$$\begin{aligned} \|\mathbf{p}_{\nu_1}(\mathbf{x}, s)\|_2 &\leq C_\sigma \cdot (\|\tau\|_2^2 + 1) \cdot \exp\left(\alpha C_\sigma \int_0^1 \int_{\mathbb{R}^k} (\|\boldsymbol{\theta}\|_2^2 + 1) d\nu_1(\boldsymbol{\theta}, s)\right) \\ &\leq C(\|\nu_1\|_\infty^2, \|\tau\|_2^2; \alpha). \end{aligned}$$

In the next, we deal with Eq. (44), define

$$\boldsymbol{\delta}_2(\mathbf{x}, s) := \mathbf{p}_{\nu_1}(\mathbf{x}, s) - \mathbf{p}_{\nu_2}(\mathbf{x}, s),$$

we have (taking $s = 1$)

$$\begin{aligned} \|\boldsymbol{\delta}_2(\mathbf{x}, 1)\|_2 &= \left\| \int_{\mathbb{R}^{k\tau}} \nabla_{\mathbf{z}} h(\mathbf{Z}_{\nu_1}(\mathbf{x}, 1), \boldsymbol{\omega}) - \nabla_{\mathbf{z}} h(\mathbf{Z}_{\nu_2}(\mathbf{x}, 1), \boldsymbol{\omega}) d\tau(\boldsymbol{\omega}) \right\|_2 \\ &\leq C_\sigma (\|\tau\|_2^2 + 1) \cdot \|\mathbf{Z}_{\nu_1}(\mathbf{x}, 1) - \mathbf{Z}_{\nu_2}(\mathbf{x}, 1)\|_2 \\ &\leq C(\|\nu_1\|_\infty^2, \|\nu_2\|_\infty^2, \|\tau\|_2^2; \alpha) \cdot \mathcal{W}_2(\nu_1, \nu_2). \end{aligned}$$

The following ODE is satisfied by $\boldsymbol{\delta}_2(\mathbf{x}, s)$ by Eq. (15),

$$\frac{\partial \boldsymbol{\delta}_2^\top(\mathbf{x}, s)}{\partial s} = -\alpha \cdot \boldsymbol{\delta}_2^\top(\mathbf{x}, \boldsymbol{\omega}, s) \int_{\mathbb{R}^k} \nabla_{\mathbf{z}} \sigma(\mathbf{Z}_{\nu_1}(\mathbf{x}, s), \boldsymbol{\theta}) d\nu_1(\boldsymbol{\theta}, s) + \alpha \cdot \mathbf{p}_{\nu_2}(\mathbf{x}, s)^\top \mathbf{D}_{\nu_1, \nu_2}(\mathbf{x}, s),$$

with

$$\mathbf{D}_{\nu_1, \nu_2}(\mathbf{x}, s) := \int_{\mathbb{R}^k} \nabla_{\mathbf{z}} \sigma(\mathbf{Z}_{\nu_2}(\mathbf{x}, s), \boldsymbol{\theta}) d\nu_2(\boldsymbol{\theta}, s) - \int_{\mathbb{R}^k} \nabla_{\mathbf{z}} \sigma(\mathbf{Z}_{\nu_1}(\mathbf{x}, s), \boldsymbol{\theta}) d\nu_1(\boldsymbol{\theta}, s).$$

Furthermore, we also split $\mathbf{D}_{\nu_1, \nu_2}(\mathbf{x}, s)$ as

$$\begin{aligned} \|\mathbf{D}_{\nu_1, \nu_2}(\mathbf{x}, s)\|_F &\leq \underbrace{\left\| \int_{\mathbb{R}^k} \nabla_{\mathbf{z}} \sigma(\mathbf{Z}_{\nu_2}(\mathbf{x}, s), \boldsymbol{\theta}) d\nu_2(\boldsymbol{\theta}, s) - \int_{\mathbb{R}^k} \nabla_{\mathbf{z}} \sigma(\mathbf{Z}_{\nu_2}(\mathbf{x}, s), \boldsymbol{\theta}) d\nu_1(\boldsymbol{\theta}, s) \right\|_F}_{(A)} \\ &\quad + \underbrace{\left\| \int_{\mathbb{R}^k} (\nabla_{\mathbf{z}} \sigma(\mathbf{Z}_{\nu_2}(\mathbf{x}, s), \boldsymbol{\theta}) - \nabla_{\mathbf{z}} \sigma(\mathbf{Z}_{\nu_1}(\mathbf{x}, s), \boldsymbol{\theta})) d\nu_1(\boldsymbol{\theta}, s) \right\|_F}_{(B)}. \end{aligned}$$

Clearly, (B) can be estimated by

$$(B) \leq C_{\mathbf{Z}}(\|\nu_1\|_{\infty}^2, \|\nu_2\|_{\infty}^2; \alpha) \cdot C_{\sigma} \cdot (\|\nu_1\|_{\infty}^2 + 1) \cdot \mathcal{W}_2(\nu_1, \nu_2).$$

To estimate (A), denote $\pi_{\nu}^* \in \Pi(\nu_1^s, \nu_2^s)$ such that $\mathbb{E}_{(\theta_1, \theta_2) \sim \pi_{\nu}^*} \|\theta_1 - \theta_2\|_2^2 = \mathcal{W}_2^2(\nu_1^s, \nu_2^s)$, by Lemma B.5, we then have

$$\begin{aligned} (A) &\leq \mathbb{E}_{(\theta_1, \theta_2) \sim \pi_{\nu}^*} \|\nabla_{\mathbf{z}} \sigma(\mathbf{Z}_{\nu_2}(\mathbf{x}, s), \theta_2) - \nabla_{\mathbf{z}} \sigma(\mathbf{Z}_{\nu_2}(\mathbf{x}, s), \theta_1)\|_F \\ &\leq C_{\sigma} \cdot \sqrt{3\mathbb{E}_{(\theta_1, \theta_2) \sim \pi_{\nu}^*} \|\theta_1\|_2^2 + \|\theta_2\|_2^2 + 1} \cdot \sqrt{\mathbb{E}_{(\theta_1, \theta_2) \sim \pi_{\nu}^*} \|\theta_1 - \theta_2\|_2^2} \\ &\leq C_{\sigma} \cdot \sqrt{3(\|\nu_1\|_{\infty}^2 + \|\nu_2\|_{\infty}^2 + 1)} \cdot \mathcal{W}_2(\nu_1, \nu_2). \end{aligned}$$

Combining the estimate of (A) and (B), we have

$$\|\mathbf{D}_{\nu_1, \nu_2}(\mathbf{x}, s)\|_F \leq C(\|\nu_1\|_{\infty}^2, \|\nu_2\|_{\infty}^2; \alpha) \cdot \mathcal{W}_2(\nu_1, \nu_2).$$

Accordingly, we are ready to estimate $\delta_2(\mathbf{x}, s)$.

$$\begin{aligned} &\frac{d\|\delta_2(\mathbf{x}, s)\|_2^2}{ds} \\ &= 2 \left(-\alpha \cdot \delta^\top(\mathbf{x}, s) \int_{\mathbb{R}^k} \nabla_{\mathbf{z}} \sigma(\mathbf{Z}_{\nu_1}(\mathbf{x}, s), \theta) d\nu_1(\theta, s) + \alpha \cdot \mathbf{p}_{\nu_2}(\mathbf{x}, s)^\top \mathbf{D}_{\nu_1, \nu_2}(\mathbf{x}, s) \right) \delta_2(\mathbf{x}, s) \\ &\leq 2\alpha \left(\|\delta_2(\mathbf{x}, s)\|_2^2 \left(1 + \int_{\mathbb{R}^k} \|\nabla_{\mathbf{z}} \sigma(\mathbf{Z}_{\nu_1}(\mathbf{x}, s), \theta)\|_F d\nu_1(\theta, s) \right) + \|\mathbf{p}_{\nu_2}(\mathbf{x}, s)^\top \mathbf{D}_{\nu_1, \nu_2}(\mathbf{x}, s)\|_2^2 \right) \\ &\leq \alpha \|\delta_2(\mathbf{x}, s)\|_2^2 (1 + C_{\sigma} \cdot (1 + \|\nu_1\|_{\infty}^2)) + 2\alpha [C(\|\nu_1\|_{\infty}^2, \|\nu_2\|_{\infty}^2, \|\tau\|_2^2; \alpha)] \cdot \mathcal{W}_2(\nu_1, \nu_2)^2 \\ &\leq C(\|\nu_1\|_{\infty}^2, \|\nu_2\|_{\infty}^2, \|\tau\|_2^2; \alpha) \cdot (\|\delta_2(\mathbf{x}, s)\|_2^2 + \mathcal{W}_2(\nu_1, \nu_2)^2). \end{aligned}$$

By the Grönwall's inequality, $\|\delta_2(\mathbf{x}, s)\|_2 \leq C_{\mathbf{p}}(\|\nu_1\|_{\infty}^2, \|\nu_2\|_{\infty}^2, \|\tau\|_2^2; \alpha) \cdot \mathcal{W}_2(\nu_1, \nu_2)$, and the proof is finished. \square

C MAIN RESULTS

C.1 GRADIENT FLOW

Proof of Theorem 4.1. To prove Theorem 4.1, we need to estimate

$$\begin{aligned} &\widehat{L}(\tau_t, \nu_t) - \widehat{L}(\tau_{t_0}, \nu_{t_0}) = \frac{1}{2} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} [(\widehat{f}_{\tau_t, \nu_t}(\mathbf{x}) - y(\mathbf{x}))^2 - (\widehat{f}_{\tau_{t_0}, \nu_{t_0}}(\mathbf{x}) - y(\mathbf{x}))^2] \\ &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} (\widehat{f}_{\tau_{t_0}, \nu_{t_0}}(\mathbf{x}) - y(\mathbf{x})) (\widehat{f}_{\tau_t, \nu_t}(\mathbf{x}) - \widehat{f}_{\tau_{t_0}, \nu_{t_0}}(\mathbf{x})) + o(|\widehat{f}_{\tau_t, \nu_t}(\mathbf{x}) - \widehat{f}_{\tau_{t_0}, \nu_{t_0}}(\mathbf{x})|), \end{aligned}$$

by $(a + \epsilon)^2 - a^2 = 2a\epsilon + o(|\epsilon|)$, where $o(\cdot)$ denotes the higher order of the error term. Then, we estimate $\widehat{f}_{\tau_t, \nu_t}(\mathbf{x}) - \widehat{f}_{\tau_{t_0}, \nu_{t_0}}(\mathbf{x})$,

$$\begin{aligned} &\widehat{f}_{\tau_t, \nu_t}(\mathbf{x}) - \widehat{f}_{\tau_{t_0}, \nu_{t_0}}(\mathbf{x}) = \beta \cdot \int_{\mathbb{R}^{k_{\tau}}} h(\mathbf{Z}_{\nu_t}(\mathbf{x}, 1), \omega) d\tau_t(\omega) - h(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, 1), \omega) d\tau_{t_0}(\omega) \\ &= \beta \cdot \left(\underbrace{\int_{\mathbb{R}^{k_{\tau}}} h(\mathbf{Z}_{\nu_t}(\mathbf{x}, 1), \omega) (d\tau_t(\omega) - d\tau_{t_0}(\omega))}_{(A)} + \underbrace{\int_{\mathbb{R}^{k_{\tau}}} (h(\mathbf{Z}_{\nu_t}(\mathbf{x}, 1), \omega) - h(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, 1), \omega)) d\tau_{t_0}(\omega)}_{(B)} \right). \end{aligned}$$

We estimate $\mathbf{Z}_{\nu_t}(\mathbf{x}, s)$ in the following, in which we assume $\boldsymbol{\theta}_t^s \sim \nu_t(\cdot, s)$, $\boldsymbol{\theta}_{t_0}^s \sim \nu_{t_0}(\cdot, s)$ in the expectation. Similar to the derivation in Lu et al. (2020); Ding et al. (2022), we have

$$\begin{aligned}
& \frac{1}{\alpha} \cdot \frac{d(\mathbf{Z}_{\nu_t} - \mathbf{Z}_{\nu_{t_0}})(\mathbf{x}, s)}{ds} = \mathbb{E}(\boldsymbol{\sigma}(\mathbf{Z}_{\nu_t}(\mathbf{x}, s), \boldsymbol{\theta}_t^s) - \boldsymbol{\sigma}(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, s), \boldsymbol{\theta}_{t_0}^s)) \\
&= \mathbb{E}(\boldsymbol{\sigma}(\mathbf{Z}_{\nu_t}(\mathbf{x}, s), \boldsymbol{\theta}_t^s) - \boldsymbol{\sigma}(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, s), \boldsymbol{\theta}_t^s)) + \mathbb{E}(\boldsymbol{\sigma}(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, s), \boldsymbol{\theta}_t^s) - \boldsymbol{\sigma}(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, s), \boldsymbol{\theta}_{t_0}^s)) \\
&= \mathbb{E} \partial_{\mathbf{z}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, s), \boldsymbol{\theta}_t^s)(\mathbf{Z}_{\nu_t}(\mathbf{x}, s) - \mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, s)) + \mathbb{E} \partial_{\boldsymbol{\theta}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, s), \boldsymbol{\theta}_t^s)(\boldsymbol{\theta}_t^s - \boldsymbol{\theta}_{t_0}^s) + o(|t - t_0|) \\
&= \mathbb{E} \partial_{\mathbf{z}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, s), \boldsymbol{\theta}_{t_0}^s)(\mathbf{Z}_{\nu_t}(\mathbf{x}, s) - \mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, s)) + \mathbb{E} \partial_{\boldsymbol{\theta}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, s), \boldsymbol{\theta}_{t_0}^s)(\boldsymbol{\theta}_t^s - \boldsymbol{\theta}_{t_0}^s) + o(|t - t_0|) \\
&= \mathbb{E} \partial_{\mathbf{z}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, s), \boldsymbol{\theta}_{t_0}^s)(\mathbf{Z}_{\nu_t}(\mathbf{x}, s) - \mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, s)) \\
&- \mathbb{E} \partial_{\boldsymbol{\theta}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, s), \boldsymbol{\theta}_{t_0}^s) \nabla_{\boldsymbol{\theta}} \frac{\delta \widehat{\mathcal{L}}(\tau_{t_0}, \nu_{t_0})}{d\nu_{t_0}}(\boldsymbol{\theta}_{t_0}^s, s)(t - t_0) + o(|t - t_0|)
\end{aligned}$$

We therefore have, by the definition of \mathbf{q}_{ν} in Eq. (17),

$$\begin{aligned}
& (\mathbf{Z}_{\nu_t} - \mathbf{Z}_{\nu_{t_0}})(\mathbf{x}, 1) \\
&= - \int_0^1 \mathbf{q}_{\nu_{t_0}}(\mathbf{x}, s) \cdot \mathbb{E} \left(\alpha \partial_{\boldsymbol{\theta}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, s), \boldsymbol{\theta}_{t_0}^s) \nabla_{\boldsymbol{\theta}} \frac{\delta \widehat{\mathcal{L}}(\tau_{t_0}, \nu_{t_0})}{d\nu_{t_0}}(\boldsymbol{\theta}, s) \right) \cdot (t - t_0) ds + o(|t - t_0|)
\end{aligned}$$

and then we have $\|(\mathbf{Z}_{\nu_t} - \mathbf{Z}_{\nu_{t_0}})(\mathbf{x}, 1)\|_2 = O(|t - t_0|)$. Using this fact and by the evolution of τ_t in Eq. (10), we estimate (A),

$$\begin{aligned}
\text{(A)} &= \int_{\mathbb{R}^{k_{\tau}}} h(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, 1), \boldsymbol{\omega})(d\tau_t(\boldsymbol{\omega}) - d\tau_{t_0}(\boldsymbol{\omega})) \\
&+ \int_{\mathbb{R}^{k_{\tau}}} (h(\mathbf{Z}_{\nu_t}(\mathbf{x}, 1), \boldsymbol{\omega}) - h(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, 1), \boldsymbol{\omega}))(d\tau_t(\boldsymbol{\omega}) - d\tau_{t_0}(\boldsymbol{\omega})) \\
&= - \int_{\mathbb{R}^{k_{\tau}}} h(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, 1), \boldsymbol{\omega}) \nabla_{\boldsymbol{\omega}} \frac{\delta \widehat{\mathcal{L}}(\tau_{t_0}, \nu_{t_0})}{\delta \tau_{t_0}}(\boldsymbol{\omega})(t - t_0) d\tau_{t_0}(\boldsymbol{\omega}) + o(|t - t_0|)
\end{aligned}$$

We can also estimate (B), in which h is hidden in the definition of \mathbf{p}_{ν} in Eq. (17),

$$\begin{aligned}
\text{(B)} &= \int_{\mathbb{R}^{k_{\tau}}} \mathbf{p}_{\nu_{t_0}}(\mathbf{x}, 1)^{\top} (\mathbf{Z}_{\nu_t} - \mathbf{Z}_{\nu_{t_0}})(\mathbf{x}, 1) d\tau_{t_0}(\boldsymbol{\omega}) + o(|t - t_0|) \\
&= - \int_{\mathbb{R}^{k_{\tau}}} \mathbf{p}_{\nu_{t_0}}(\mathbf{x}, s)^{\top} \\
&\quad \left(\mathbb{E} \alpha \nabla_{\boldsymbol{\theta}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, s), \boldsymbol{\theta}_{t_0}^s) \nabla_{\boldsymbol{\theta}} \frac{\delta \widehat{\mathcal{L}}(\tau_{t_0}, \nu_{t_0})}{d\nu_{t_0}}(\boldsymbol{\theta}_{t_0}^s, s) \cdot (t - t_0) \right) d\tau_{t_0}(\boldsymbol{\omega}) + o(|t - t_0|) \\
&= - \int_{\mathbb{R}^{k_{\tau}} \times \mathbb{R}^{k_{\nu}} \times [0, 1]} \left(\mathbf{p}_{\nu_{t_0}}(\mathbf{x}, s)^{\top} \cdot \alpha \nabla_{\boldsymbol{\theta}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, s), \boldsymbol{\theta}) \cdot \right. \\
&\quad \left. \nabla_{\boldsymbol{\theta}} \frac{\delta \widehat{\mathcal{L}}(\tau_{t_0}, \nu_{t_0})}{\delta \nu_{t_0}}(\boldsymbol{\theta}, s) \right) d\nu_{t_0}(\boldsymbol{\theta}, s) \cdot (t - t_0) + o(|t - t_0|).
\end{aligned}$$

Combine the estimation of (A) and (B), we have

$$\begin{aligned}
& \widehat{\mathcal{L}}(\tau_t, \nu_t) - \widehat{\mathcal{L}}(\tau_{t_0}, \nu_{t_0}) = \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} \beta(\widehat{f}_{\tau_{t_0}, \nu_{t_0}}(\mathbf{x}) - y(\mathbf{x}))((\text{A}) + (\text{B})) \\
&= - \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} \beta(\widehat{f}_{\tau_{t_0}, \nu_{t_0}}(\mathbf{x}) - y(\mathbf{x})) \int_{\mathbb{R}^{k_{\tau}} \times \mathbb{R}^k \times [0, 1]} d\tau_{t_0}(\boldsymbol{\omega}) d\nu_{t_0}(\boldsymbol{\theta}, s)(t - t_0) \\
&\quad \cdot \left(\mathbf{Z}_{\nu_{t_0}}^{\top}(\mathbf{x}, 1) \nabla_{\boldsymbol{\omega}} \frac{\delta \widehat{\mathcal{L}}(\tau_{t_0}, \nu_{t_0})}{\delta \tau_{t_0}}(\boldsymbol{\omega}) \right. \\
&\quad \left. + \mathbf{p}_{\nu_{t_0}}^{\top}(\mathbf{x}, s) \cdot \alpha \nabla_{\boldsymbol{\theta}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_{t_0}}(\mathbf{x}, s), \boldsymbol{\theta}) \nabla_{\boldsymbol{\theta}} \frac{\delta \widehat{\mathcal{L}}(\tau_{t_0}, \nu_{t_0})}{\delta \nu_{t_0}}(\boldsymbol{\theta}, s) + o(|t - t_0|) \right) \\
&= - \mathbb{E}_{\boldsymbol{\omega} \sim \tau_{t_0}, (\boldsymbol{\theta}, s) \sim \nu_{t_0}} \left(\left\| \nabla_{\boldsymbol{\theta}} \frac{\delta \widehat{\mathcal{L}}(\tau, \nu)}{\delta \nu}(\boldsymbol{\theta}, s) \right\|_2^2 + \left\| \nabla_{\boldsymbol{\omega}} \frac{\delta \widehat{\mathcal{L}}(\tau, \nu)}{\delta \tau}(\boldsymbol{\omega}) \right\|_2^2 \right) (t - t_0) + o(|t - t_0|),
\end{aligned}$$

from the definition of functional gradient in Eq. (11) and Eq. (14). In all, the theorem is proved. \square

Proof of Proposition 4.2. We expand the function derivative:

$$\begin{aligned}
& \int_{\mathbb{R}^{k_\tau} \times \mathbb{R}^{k_\nu} \times [0,1]} \left\| \nabla_{\boldsymbol{\theta}} \frac{\delta \widehat{L}(\tau_t, \nu_t)}{\delta \nu_t}(\boldsymbol{\theta}, s) \right\|_2^2 d\tau_t(\boldsymbol{\omega}) d\nu_t(\boldsymbol{\theta}, s) \\
&= \int_{\mathbb{R}^{k_\tau} \times \mathbb{R}^{k_\nu} \times [0,1]} \frac{\beta^2 \alpha^2}{n^2} \sum_{i,j=1}^n (\widehat{f}_{\tau_t, \nu_t}(\mathbf{x}_i) - y(\mathbf{x}_i)) (\widehat{f}_{\tau_t, \nu_t}(\mathbf{x}_j) - y(\mathbf{x}_j)) \\
&\quad \cdot \mathbf{p}_{\nu_t}^\top(\mathbf{x}_i, s) \nabla_{\boldsymbol{\theta}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_t}(\mathbf{x}_i, s), \boldsymbol{\theta}) \nabla_{\boldsymbol{\theta}}^\top \boldsymbol{\sigma}(\mathbf{Z}_{\nu_t}(\mathbf{x}_j, s), \boldsymbol{\theta}) \mathbf{p}_{\nu_t}(\mathbf{x}_j, s) d\tau_t(\boldsymbol{\omega}) d\nu_t(\boldsymbol{\theta}, s) \\
&= \frac{\alpha^2 \beta^2}{n^2} \mathbf{b}_t^\top \mathbf{G}_1(\tau_t, \nu_t) \mathbf{b}_t,
\end{aligned}$$

and similarly,

$$\int_{\mathbb{R}^{k_\tau}} \left\| \nabla_{\boldsymbol{\omega}} \frac{\delta \widehat{L}(\tau_t, \nu_t)}{\delta \tau_t}(\boldsymbol{\omega}) \right\|_2^2 d\tau_t(\boldsymbol{\omega}) = \frac{\beta^2}{n^2} \mathbf{b}_t^\top \mathbf{G}_2(\tau_t, \nu_t) \mathbf{b}_t.$$

In all, the lemma is proved. \square

Proof of Lemma 4.6. We use the expansion of the gradient flow:

$$\begin{aligned}
\frac{d\text{KL}(\tau_t \| \tau_0)}{dt} &= \int_{\mathbb{R}^{k_\tau}} \frac{\delta \text{KL}(\tau_t \| \tau_0)}{\delta \tau_t} \frac{d\tau_t}{dt} d\boldsymbol{\omega} = \int_{\mathbb{R}^{k_\tau}} \frac{\delta \text{KL}(\tau_t \| \tau_0)}{\delta \tau_t} \nabla \cdot \left(\tau_t(\boldsymbol{\omega}) \nabla \frac{\delta \widehat{L}(\tau_t, \nu_t)}{\delta \tau_t} \right) d\boldsymbol{\omega} \\
&= - \int_{\mathbb{R}^{k_\tau}} \tau_t(\boldsymbol{\omega}) \left(\nabla \frac{\delta \text{KL}(\tau_t \| \tau_0)}{\delta \tau_t} \right) \left(\nabla \frac{\delta \widehat{L}(\tau_t, \nu_t)}{\delta \tau_t} \right) d\boldsymbol{\omega}.
\end{aligned}$$

Similarly, we have

$$\begin{aligned}
\frac{d\text{KL}(\nu_t^s \| \nu_0^s)}{dt} &= \int_{\mathbb{R}^{k_\nu}} \frac{\delta \text{KL}(\nu_t^s \| \nu_0^s)}{\delta \nu_t^s} \frac{d\nu_t^s}{dt} d\boldsymbol{\theta} \\
&= \int_{\mathbb{R}^{k_\nu}} \frac{\delta \text{KL}(\nu_t^s \| \nu_0^s)}{\delta \nu_t^s} \nabla_{\boldsymbol{\theta}} \cdot \left(\nu_t^s(\boldsymbol{\theta}) \nabla_{\boldsymbol{\theta}} \frac{\delta \widehat{L}(\tau_t, \nu_t)}{\delta \nu_t}(\boldsymbol{\theta}, s) \right) d\boldsymbol{\theta} \\
&= - \int_{\mathbb{R}^{k_\nu}} \nu_t^s(\boldsymbol{\theta}) \left(\nabla_{\boldsymbol{\theta}} \frac{\delta \text{KL}(\nu_t^s \| \nu_0^s)}{\delta \nu_t^s} \right) \left(\nabla_{\boldsymbol{\theta}} \frac{\delta \widehat{L}(\tau_t, \nu_t)}{\delta \nu_t}(\boldsymbol{\theta}, s) \right) d\boldsymbol{\theta}.
\end{aligned}$$

Therefore, the proof is completed. \square

C.2 MINIMUM EIGENVALUE AT INITIALIZATION

Proof of Lemma 4.3. In the proof, similar to Assumption 3.3, we assume $\boldsymbol{\theta} = (\mathbf{u}, \mathbf{w}, b) \in \mathbb{R}^{2d+1}$, $\boldsymbol{\omega} = (a, \mathbf{w}, b) \in \mathbb{R}^{d+2}$, $\mathbf{u}, \mathbf{w} \in \mathbb{R}^d$, $a, b \in \mathbb{R}$. At initialization, we notice that $\nu_0(\boldsymbol{\theta}, s) \propto \exp\left(-\frac{\|\boldsymbol{\theta}\|_2^2}{2}\right)$, and $\tau_0(\boldsymbol{\omega}) \propto \exp\left(-\frac{\|\boldsymbol{\omega}\|_2^2}{2}\right)$ are standard Gaussian. Since the distribution of \mathbf{u} , a is symmetric, and independent from other parts of $\boldsymbol{\theta}$, $\boldsymbol{\omega}$ respectively, we have

$$\begin{aligned}
\frac{d\mathbf{Z}_{\nu_0}(\mathbf{x}, s)}{ds} &= \int_{\mathbb{R}^{k_\nu}} \mathbf{u}^\top \sigma_0(\mathbf{w}^\top \mathbf{Z}_{\nu_0}(\mathbf{x}, s) + b) d\nu_0(\mathbf{u}, \mathbf{w}, b, s) \\
&= \int_{\mathbb{R}^d} \mathbf{u}^\top d\nu_0(\mathbf{u}) \int_{\mathbb{R}^{k_\nu-d}} \sigma_0(\mathbf{w}^\top \mathbf{Z}_{\nu_0}(\mathbf{x}, s) + b) d\nu_0(\mathbf{w}, b, s) = \mathbf{0}, \forall s \in [0, 1], \\
\frac{d\mathbf{p}_{\nu_0}^\top(\mathbf{x}, s)}{ds} &= -\alpha \cdot \mathbf{p}_{\nu_0}^\top(\mathbf{x}, s) \int_{\mathbb{R}^{k_\nu}} \nabla_{\mathbf{z}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_0}(\mathbf{x}, s), \boldsymbol{\theta}) d\nu_0(\boldsymbol{\theta}, s) \\
&= -\alpha \cdot \mathbf{p}_{\nu_0}^\top(\mathbf{x}, s) \int_{\mathbb{R}^{k_\nu}} \mathbf{u} \mathbf{w}^\top \sigma_0'(\mathbf{w}^\top \mathbf{Z}_{\nu_0}(\mathbf{x}, s) + b) d\nu_0(\mathbf{u}, \mathbf{w}, b, s) = \mathbf{0}, \\
\mathbf{p}_{\nu_0}(\mathbf{x}, 1) &= \int_{\mathbb{R}^{k_\tau}} \nabla_{\mathbf{z}}^\top h(\mathbf{Z}_{\nu_0}(\mathbf{x}, 1), \boldsymbol{\omega}) d\tau_0(\boldsymbol{\omega}) \\
&= \int_{\mathbb{R}^{k_\tau}} a \mathbf{w} \sigma_0'(\mathbf{w}^\top \mathbf{Z}_{\nu_0}(\mathbf{x}, 1) + b) \tau_0(a, \mathbf{w}, b) = \mathbf{0}.
\end{aligned}$$

From the first two equations, we have

$$\mathbf{Z}_{\nu_0}(\mathbf{x}, s) = \mathbf{x}, \forall s \quad \mathbf{p}_{\nu_0}(\mathbf{x}, s) = \mathbf{p}_{\nu_0}(\mathbf{x}, 1) = 0$$

By the definition of $\mathbf{G}_2(\tau_0, \nu_0)$, we have

$$\begin{aligned} \mathbf{G}_2(\tau_0, \sigma_0) &= \mathbb{E}_{(a, \mathbf{w}, b) \sim \mathcal{N}(0, I)} \nabla_{\omega} h(\mathbf{X}, \omega) \nabla_{\omega}^{\top} h(\mathbf{X}, \omega) \\ &= \mathbb{E}_{(a, \mathbf{w}, b) \sim \mathcal{N}(0, I)} (\boldsymbol{\sigma}_0((\mathbf{X}, \mathbb{1})(\mathbf{w}, b)), a\boldsymbol{\sigma}'_0((\mathbf{X}, \mathbb{1})(\mathbf{w}, b)), \boldsymbol{\sigma}'_0((\mathbf{X}, \mathbb{1})(\mathbf{w}, b))) \\ &\quad (\boldsymbol{\sigma}_0((\mathbf{X}, \mathbb{1})(\mathbf{w}, b)), a\boldsymbol{\sigma}'_0((\mathbf{X}, \mathbb{1})(\mathbf{w}, b)), \boldsymbol{\sigma}'_0((\mathbf{X}, \mathbb{1})(\mathbf{w}, b)))^{\top}, \\ &\geq \mathbb{E}_{(a, \mathbf{w}, b) \sim \mathcal{N}(0, I)} \boldsymbol{\sigma}_0((\mathbf{X}, \mathbb{1})(\mathbf{w}, b)) \boldsymbol{\sigma}_0((\mathbf{X}, \mathbb{1})(\mathbf{w}, b))^{\top}. \end{aligned}$$

Let $\bar{\mathbf{x}} = (\mathbf{x}, 1)$, by Assumption 3.1, the cosine similarity of $\bar{\mathbf{x}}_i$ and $\bar{\mathbf{x}}_j$ is no larger than $(1 + C_{\max})/2$. Then we bound $\lambda_{\min}(\mathbf{G}^{(2)})$:

$$\begin{aligned} \lambda_{\min}(\mathbf{G}^{(2)}) &\geq \lambda_{\min}\left(\mathbb{E}_{\mathbf{w} \sim \mathcal{N}(0, \mathbb{I}_{d+1})} [\sigma_1(\bar{\mathbf{X}}\mathbf{w})\sigma_1(\bar{\mathbf{X}}\mathbf{w})^{\top}]\right) \\ &= \lambda_{\min}\left(\sum_{s=0}^{\infty} \mu_s(\sigma_1)^2 \circ_{i=1}^s (\bar{\mathbf{X}}\bar{\mathbf{X}}^{\top})\right) \quad (\text{Nguyen \& Mondelli, 2020, Lemma D.3}) \\ &\geq \mu_r(\sigma_1)^2 \lambda_{\min}(\circ_{i=1}^r \bar{\mathbf{X}}\bar{\mathbf{X}}^{\top}) \quad \left(\text{taking } r \geq \frac{2 \log(2n)}{1 - C_{\max}}\right) \\ &\geq \mu_r(\sigma_1)^2 \left(\min_{i \in [n]} \|\bar{\mathbf{x}}_i\|_2^{2r} - (n-1) \max_{i \neq j} |\langle \bar{\mathbf{x}}_i, \bar{\mathbf{x}}_j \rangle|^r\right) \quad [\text{Gershgorin circle theorem}] \\ &\geq \mu_r(\sigma_1)^2 \left(1 - (n-1) \left(\frac{1 + C_{\max}}{2}\right)^r\right), \\ &\geq \mu_r(\sigma_1)^2 \left(1 - (n-1) \left(1 - \frac{\log(2n)}{r}\right)^r\right) \\ &\geq \mu_r(\sigma_1)^2 \left(1 - (n-1) \exp(-\log(2n))\right) \\ &\geq \mu_r(\sigma_1)^2 / 2, \end{aligned}$$

where the last inequality holds by the fact that $(1 - \frac{\log(2n)}{r})^r$ is an increasing function of r . \square

C.3 PERTURBATION OF MINIMUM EIGENVALUE

In this section, we analyze the minimum eigenvalue of the Gram matrix.

Lemma C.1. *The perturbation of $\mathbf{G}_2(\tau, \nu)$ can be upper bounded in the following, for any $i, j \in [n]$,*

$$|\mathbf{G}_2(\tau, \nu) - \mathbf{G}_2(\tau_0, \nu_0)|_{i,j} \leq C_{\mathbf{G}}(\|\tau\|_2^2, \|\nu\|_{\infty}^2; d, \alpha)(\mathcal{W}_2(\tau, \tau_0) + \mathcal{W}_2(\nu, \nu_0)),$$

where \mathbf{G}_2 is defined in Section 4.1, and τ_0, ν_0 satisfies Assumption 3.2.

Proof of Lemma C.1. We deal with $\mathbf{G}_2(\tau, \nu)$ in an element-wise way. Let $(\omega, \omega_0) \sim \pi_{\tau}^*$ be the optimal coupling of $\mathcal{W}_2(\tau, \tau_0)$, the difference can be estimated by

$$\begin{aligned} &|\mathbf{G}_2(\tau, \nu) - \mathbf{G}_2(\tau_0, \nu_0)|_{i,j} \\ &\leq \mathbb{E} |\nabla_{\omega} h(\mathbf{Z}_{\nu}(\mathbf{x}_i, 1), \omega) \nabla_{\omega}^{\top} h(\mathbf{Z}_{\nu}(\mathbf{x}_j, 1), \omega) - \nabla_{\omega} h(\mathbf{Z}_{\nu_0}(\mathbf{x}_i, 1), \omega_0) \nabla_{\omega}^{\top} h(\mathbf{Z}_{\nu_0}(\mathbf{x}_j, 1), \omega_0)| \\ &\leq \underbrace{\mathbb{E} |\nabla_{\omega} h(\mathbf{Z}_{\nu}(\mathbf{x}_i, 1), \omega) (\nabla_{\omega} h(\mathbf{Z}_{\nu}(\mathbf{x}_j, 1), \omega) - \nabla_{\omega} h(\mathbf{Z}_{\nu_0}(\mathbf{x}_j, 1), \omega_0))^{\top}|}_{(A)} \\ &\quad + \underbrace{\mathbb{E} |(\nabla_{\omega} h(\mathbf{Z}_{\nu}(\mathbf{x}_i, 1), \omega) - \nabla_{\omega} h(\mathbf{Z}_{\nu_0}(\mathbf{x}_i, 1), \omega_0)) \nabla_{\omega}^{\top} h(\mathbf{Z}_{\nu_0}(\mathbf{x}_j, 1), \omega_0)|}_{(B)}. \end{aligned}$$

We then estimate (A) and (B) separately. The term (A) involves

$$\begin{aligned} & \|\nabla_{\omega} h(\mathbf{Z}_{\nu}(\mathbf{x}, 1), \omega) - \nabla_{\omega} h(\mathbf{Z}_{\nu_0}(\mathbf{x}, 1), \omega_0)\|_2 \\ & \leq \|\nabla_{\omega} h(\mathbf{Z}_{\nu}(\mathbf{x}, 1), \omega) - \nabla_{\omega} h(\mathbf{Z}_{\nu}(\mathbf{x}, 1), \omega_0)\|_2 + \|\nabla_{\omega} h(\mathbf{Z}_{\nu}(\mathbf{x}, 1), \omega_0) - \nabla_{\omega} h(\mathbf{Z}_{\nu_0}(\mathbf{x}, 1), \omega_0)\|_2 \\ & \leq C_{\sigma} \cdot (\|\mathbf{Z}_{\nu}(\mathbf{x}, 1)\|_2 + 1) \cdot \|\omega - \omega_0\|_2 + C_{\sigma} \cdot (\|\omega_0\|_2^2 + 1) \cdot \|\mathbf{Z}_{\nu}(\mathbf{x}, 1) - \mathbf{Z}_{\nu_0}(\mathbf{x}, 1)\|_2 \\ & \leq C_{\sigma} \cdot ((C_{\mathbf{Z}}(\|\nu\|_{\infty}^2; \alpha) + 1) \cdot \|\omega - \omega_0\|_2 + (\|\omega_0\|_2^2 + 1) \cdot C_{\mathbf{Z}}(\|\nu\|_{\infty}^2, \|\nu_0\|_{\infty}^2; \alpha) \mathcal{W}_2(\nu, \nu_0)), \end{aligned}$$

where we use Lemmas B.5 and B.6 in our proof. Besides, the term (B) involves

$$\begin{aligned} \|\nabla_{\omega} h(\mathbf{Z}_{\nu}(\mathbf{x}, 1), \omega)\|_2 & \leq C_{\sigma} \cdot (\|\mathbf{Z}_{\nu}(\mathbf{x}, 1)\|_2 + 1) \cdot (\|\omega\|_2 + 1) \\ & \leq C_{\sigma} \cdot (C_{\mathbf{Z}}(\|\nu\|_{\infty}^2; \alpha) + 1) \cdot (\|\omega\|_2 + 1). \end{aligned}$$

Therefore, by $\mathbb{E}\|\omega_0\|_2^4 = 3k_{\tau} = 3(d+2)$.

$$\begin{aligned} \text{(A) + (B)} & \leq C(\|\nu\|_{\infty}^2, \|\nu_0\|_{\infty}^2; \alpha) \mathbb{E}(\|\omega - \omega_0\|_2 + (\|\omega_0\|_2^2 + 1) \mathcal{W}_2(\nu, \nu_0)) (\|\omega\|_2 + \|\omega_0\|_2 + 2) \\ & \leq C(\|\nu\|_{\infty}^2, \|\nu_0\|_{\infty}^2; \alpha) \mathbb{E}(\|\omega - \omega_0\|_2 + (\|\omega_0\|_2^2 + 1) \mathcal{W}_2(\nu, \nu_0)) (\|\omega\|_2^2 + \|\omega_0\|_2^2 + 4) \\ & \leq C(\|\tau\|_2^2, \|\tau_0\|_2^2, \|\nu\|_{\infty}^2, \|\nu_0\|_{\infty}^2; d, \alpha) (\mathbb{E}\|\omega - \omega_0\|_2 (\|\omega\|_2^2 + \|\omega_0\|_2^2 + 4) + \mathcal{W}_2(\nu, \nu_0)) \\ & \leq C(\|\tau\|_2^2, \|\tau_0\|_2^2, \|\nu\|_{\infty}^2, \|\nu_0\|_{\infty}^2; d, \alpha) [(\mathbb{E}\|\omega - \omega_0\|_2^2)^{\frac{1}{2}} (\mathbb{E}(\|\omega\|_2^2 + \|\omega_0\|_2^2 + 4)^2)^{\frac{1}{2}} + \mathcal{W}_2(\nu, \nu_0)] \\ & \leq C(\|\tau\|_2^2, \|\tau_0\|_2^2, \|\nu\|_{\infty}^2, \|\nu_0\|_{\infty}^2; d, \alpha) (\mathcal{W}_2(\tau, \tau_0) + \mathcal{W}_2(\nu, \nu_0)), \end{aligned}$$

since $\mathbb{E}\|\omega - \omega_0\|_2^2 = (\mathcal{W}_2(\tau, \tau_0))^2$, by the definition of optimal coupling. Since $\|\tau_0\|_2^2 = d + 2$, $\|\nu_0\|_{\infty}^2 = 2d + 1$, we can drop dependence of C on $\|\tau_0\|_2^2$, $\|\nu_0\|_{\infty}^2$ and replace them by d . In all, the lemma is proved. Specifically, we could set

$$\begin{aligned} & C_{\mathbf{G}}(\|\tau\|_2^2, \|\nu\|_{\infty}^2; d, \alpha) \\ & := 16(d+1)C_{\sigma}^2(C_{\mathbf{Z}}(\|\nu\|_{\infty}^2; \alpha) + 1) + C_{\mathbf{Z}}(\|\nu\|_{\infty}^2, 2d+1; \alpha)^2 (\|\tau\|_2^2 + d + 1) \end{aligned}$$

□

Lemma C.2. *If $\nu \in \mathcal{C}(\mathcal{P}^2; [0, 1])$, $\tau \in \mathcal{P}^2$, $\mathcal{W}_2(\nu, \nu_0) \leq \sqrt{d}$, and $\mathcal{W}_2(\tau, \tau_0) \leq \sqrt{d}$, we have $\forall i, j \in [n]$,*

$$|\mathbf{G}_2(\tau, \nu) - \mathbf{G}_2(\tau_0, \nu_0)|_{i,j} \leq C_{\mathbf{G}}(d, \alpha) \cdot (\mathcal{W}_2(\nu, \nu_0) + \mathcal{W}_2(\tau, \tau_0))$$

Proof of Lemma C.2. For any $s \in [0, 1]$, let $(\theta^s, \theta_0^s) \sim \pi_{\nu^s}^*$ be the optimal coupling of $\mathcal{W}_2(\nu^s, \nu_0^s)$.

$$\|\nu^s\|_2^2 = \mathbb{E}\|\theta^s\|_2^2 \leq 2\mathbb{E}(\|\theta^s - \theta_0^s\|_2^2 + \|\theta_0^s\|_2^2) = 2\mathcal{W}_2^2(\nu^s, \nu_0^s) + 2(2d+1) \leq 6d+2.$$

where the last inequality holds, since $\mathcal{W}_2(\nu^s, \nu_0^s) \leq \mathcal{W}_2(\nu, \nu_0)$, $\forall s \in [0, 1]$.

We also let $(\omega, \omega_0) \sim \pi_{\tau}^*$ be the optimal coupling of $\mathcal{W}_2(\tau, \tau_0)$.

$$\|\tau\|_2^2 = \mathbb{E}\|\omega\|_2^2 \leq 2\mathbb{E}(\|\omega - \omega_0\|_2^2 + \|\omega_0\|_2^2) = 2\mathcal{W}_2^2(\tau, \tau_0) + 2(d+2) \leq 4(d+1).$$

Therefore, $\|\nu\|_{\infty} \leq \sqrt{6d+2}$, $\|\tau\|_2 \leq 2\sqrt{d+1}$.

By Lemma C.1, replacing $\|\tau\|_2^2$, $\|\nu\|_{\infty}^2$ with their upper bound w.r.t. d in the definition of $C_{\mathbf{G}}(\|\tau\|_2^2, \|\nu\|_{\infty}^2; d, \alpha)$, there exist $C_{\mathbf{G}}(d, \alpha)$ satisfying Lemma C.2. □

For ease of description, we restate Lemma 4.4 with more details here.

Lemma C.3. *If $\nu \in \mathcal{C}(\mathcal{P}^2; [0, 1])$, $\tau \in \mathcal{P}^2$, $\mathcal{W}_2(\nu, \nu_0) \leq r$ and $\mathcal{W}_2(\tau, \tau_0) \leq r$, we have*

$$\lambda_{\min}(\mathbf{G}_2(\tau, \nu)) \geq \frac{\Lambda}{2}, \text{ with } r := r_{\max}(d, \alpha) = \min \left\{ \sqrt{d}, \frac{\Lambda}{4nC_{\mathbf{G}}(d, \alpha)} \right\}.$$

where Λ is defined in Lemma 4.3, and $C_{\mathbf{G}}(d, \alpha)$ is defined in Lemma C.2.

Proof of Lemma C.3. By Lemma C.2, let

$$r := \min \left\{ \sqrt{d}, \frac{\Lambda}{4nC_{\mathbf{G}}(d, \alpha)} \right\},$$

By Lemma 4.3, we have $\forall i, j \in [n]$,

$$|\mathbf{G}_2(\tau, \nu) - \mathbf{G}_2(\tau_0, \nu_0)|_{i,j} \leq C_G(d; \alpha) \cdot (\mathcal{W}_2(\nu, \nu_0) + \mathcal{W}_2(\tau, \tau_0)) \leq \frac{\Lambda}{2n}.$$

By the standard matrix perturbation bounds, we have

$$\begin{aligned} \lambda_{\min}(\mathbf{G}_2(\tau, \nu)) &\geq \lambda_{\min}(\mathbf{G}_2(\tau_0, \nu_0)) - \|\mathbf{G}_2(\tau, \nu) - \mathbf{G}_2(\tau_0, \nu_0)\|_2 \\ &\geq \lambda_{\min}(\mathbf{G}_2(\tau_0, \nu_0)) - n\|\mathbf{G}_2(\tau, \nu) - \mathbf{G}_2(\tau_0, \nu_0)\|_{\infty, \infty} \\ &\geq \frac{\Lambda}{2}. \end{aligned}$$

□

C.4 ESTIMATION OF KL DIVERGENCE.

Inspired by Lemma C.3, we propose the following definition:

Definition C.4. *Define*

$$t_{\max} := \sup\{t_0, \text{s.t. } \forall t \in [0, t_0], \max\{\mathcal{W}_2(\nu_t, \nu_0), \mathcal{W}_2(\tau_t, \tau_0)\} \leq r_{\max}\},$$

where r_{\max} is defined in Lemma 4.4.

By Lemma C.2 and Lemma C.3, for $t \leq t_{\max}$, we have $\max\{\|\nu_t\|_{\infty}^2, \|\tau_t\|_2^2\} = O(d)$, and $\lambda_{\min}(\mathbf{G}_2(\tau_t, \nu_t)) \geq \frac{\Lambda}{2}$.

We first prove the linear convergence of empirical loss under finite time.

Lemma C.5. *Assume the PDE (10) has solution $\tau_t \in \mathcal{P}^2$, and the PDE (12) has solution $\nu_t \in \mathcal{C}(\mathcal{P}^2; [0, 1])$. Under Assumption 3.1, 3.2, 3.3, for all $t \in [0, t_{\max})$, we have*

$$\widehat{L}(\tau_t, \nu_t) \leq e^{-\frac{\beta^2 \Lambda}{2n} t} \widehat{L}(\tau_0, \nu_0), \quad \text{KL}(\tau_t \| \tau_0) \leq \frac{C_{\text{KL}}(d, \alpha)}{\Lambda^2 \beta^2}, \quad \text{KL}(\nu_t \| \nu_0) \leq \frac{C_{\text{KL}}(d, \alpha)}{\Lambda^2 \beta^2}. \quad (45)$$

Proof of Lemma C.5. Please see Lemma C.6, Lemma C.7, and Lemma C.8. □

Lemma C.6. *Assume τ_t, ν_t is the solution to PDE (10) and (12), we have for $t < t_{\max}$,*

$$\widehat{L}(\tau_t, \nu_t) \leq e^{-\frac{\beta^2 \Lambda}{2n} t} \widehat{L}(\tau_0, \nu_0),$$

where Λ is defined in Lemma 4.3.

Proof of Lemma C.6. By Lemma C.3, for $t < t_{\max}$, $\lambda_{\min}(\mathbf{G}(\tau_t, \nu_t)) \geq \frac{\Lambda}{2}$,

$$\frac{\partial \widehat{L}(\nu_t, \tau_t)}{\partial t} = -\frac{\beta^2}{n^2} \mathbf{b}_t^\top (\alpha^2 \mathbf{G}_1(\tau_t, \nu_t) + \mathbf{G}_2(\tau_t, \nu_t)) \mathbf{b}_t \leq -\frac{\beta^2 \Lambda}{2n} \widehat{L}(\tau_t, \nu_t) \leq -\frac{\beta^2 \Lambda}{2n} \widehat{L}(\tau_0, \nu_0).$$

Therefore, we have

$$\widehat{L}(\tau_t, \nu_t) \leq e^{-\frac{\beta^2 \Lambda}{2n} t} \widehat{L}(\tau_0, \nu_0).$$

□

Lemma C.7. *Assume the PDE (12) has solution $\nu_t \in \mathcal{C}(\mathcal{P}^2; [0, 1])$, and the PDE (10) has solution $\tau_t \in \mathcal{P}^2$. Under Assumption 3.3, 3.1, 3.2, then for all $t \in [0, t_{\max})$, the following results hold,*

$$\text{KL}(\nu_t^s \| \nu_0^s) \leq \frac{1}{\Lambda^2 \beta^2} C_{\text{KL}}(d, \alpha), \quad \forall s \in [0, 1].$$

Proof of Lemma C.7. By Gaussian initialization of ν_0^s , $\log \nu_0^s(\boldsymbol{\theta}) = -\frac{\|\boldsymbol{\theta}\|_2^2}{2} + C$, and we thus have $\nabla_{\boldsymbol{\theta}} \frac{\partial \text{KL}(\nu_t^s \| \nu_0^s)}{\partial \nu_t^s} = \nabla_{\boldsymbol{\theta}} \log \nu_t^s + \boldsymbol{\theta}$. Combining this with Lemma 4.6 and Eq. (14), we have

$$\begin{aligned} \frac{\partial \text{KL}(\nu_t^s \| \nu_0^s)}{\partial t} &= -\mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} \beta \cdot (f_{\tau_t, \nu_t}(\mathbf{x}) - y(\mathbf{x})) \\ &\quad \cdot \alpha \int_{\mathbb{R}^{k_\tau}} \nabla_{\boldsymbol{\theta}} (\mathbf{p}_{\nu_t}^\top(\mathbf{x}, s) \boldsymbol{\sigma}(\mathbf{Z}_{\nu_t}(\mathbf{x}, s), \boldsymbol{\theta})) \cdot (\nabla_{\boldsymbol{\theta}} \log \nu_t^s + \boldsymbol{\theta}) d\nu_t^s(\boldsymbol{\theta}). \end{aligned}$$

Define

$$J_t^s(\mathbf{x}, \boldsymbol{\theta}) := -\alpha \cdot \mathbf{p}_{\nu_t}(\mathbf{x}, s)^\top \left(\nabla_{\boldsymbol{\theta}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_t}(\mathbf{x}, s), \boldsymbol{\theta}) \cdot \boldsymbol{\theta} - \Delta_{\boldsymbol{\theta}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_t}(\mathbf{x}, s), \boldsymbol{\theta}) \right).$$

By the definition of $J_t^s(\mathbf{x}, \boldsymbol{\theta})$, and integration by parts, we have

$$\frac{\partial \text{KL}(\nu_t^s \parallel \nu_0^s)}{\partial t} = \beta \cdot \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} [(f_{\tau_t, \nu_t}(\mathbf{x}) - y(\mathbf{x})) \mathbb{E}_{\boldsymbol{\theta} \sim \nu_t^s} J_t^s(\mathbf{x}, \boldsymbol{\theta})].$$

We can obtain the gradient of J_t^s w.r.t. $\boldsymbol{\theta}$,

$$\nabla_{\boldsymbol{\theta}} J_t^s(\mathbf{x}, \boldsymbol{\theta}) = -\alpha \cdot \mathbf{p}_{\nu_t}(\mathbf{x}, s)^\top \left(\nabla_{\boldsymbol{\theta}} (\nabla_{\boldsymbol{\theta}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_t}(\mathbf{x}, s), \boldsymbol{\theta}) \cdot \boldsymbol{\theta}) - \nabla_{\boldsymbol{\theta}} \Delta_{\boldsymbol{\theta}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_t}(\mathbf{x}, s), \boldsymbol{\theta}) \right).$$

Therefore, by Lemma B.4, and the estimate $\|\mathbf{Z}_{\nu_t}(\mathbf{x}, s)\|_2 \leq C_{\mathbf{Z}}(\|\nu_t\|_\infty^2; \alpha)$ from Lemma B.6, we have

$$\begin{aligned} \|\nabla_{\boldsymbol{\theta}} (\nabla_{\boldsymbol{\theta}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_t}(\mathbf{x}, s), \boldsymbol{\theta}) \cdot \boldsymbol{\theta})\|_F &\leq C_{\boldsymbol{\sigma}} \cdot (\|\boldsymbol{\theta}\|_2 + 1) \cdot (C_{\mathbf{Z}}(\|\nu_t\|_\infty^2; \alpha) + 1) \\ \|\nabla_{\boldsymbol{\theta}} \Delta_{\boldsymbol{\theta}} \boldsymbol{\sigma}(\mathbf{Z}_{\nu_t}(\mathbf{x}, s), \boldsymbol{\theta})\|_F &\leq C_{\boldsymbol{\sigma}} \cdot (\|\boldsymbol{\theta}\|_2 + 1) \cdot (C_{\mathbf{Z}}(\|\nu_t\|_\infty^2; \alpha)^3 + 1). \end{aligned}$$

Therefore, we can estimate $\nabla_{\boldsymbol{\theta}} J_t^s(\mathbf{x}, \boldsymbol{\theta})$,

$$\begin{aligned} \|\nabla_{\boldsymbol{\theta}} J_t^s(\mathbf{x}, \boldsymbol{\theta})\|_2 &\leq 2C_{\boldsymbol{\sigma}} \cdot (C_{\mathbf{Z}}(\|\nu_t\|_\infty^2; \alpha)^3 + 1) (\|\boldsymbol{\theta}\|_2 + 1) \|\mathbf{p}_{\nu_t}(\mathbf{x}, s)\|_2 \\ &\leq 2C_{\boldsymbol{\sigma}} \cdot (C_{\mathbf{Z}}(\|\nu_t\|_\infty^2; \alpha)^3 + 1) \cdot C_{\mathbf{p}}(\|\nu_t\|_\infty^2, \|\tau_t\|_2^2; \alpha) \cdot (\|\boldsymbol{\theta}\|_2 + 1) \\ &\leq C(\|\tau_t\|_2^2, \|\nu_t\|_\infty^2; \alpha) \cdot (\|\boldsymbol{\theta}\|_2 + 1). \end{aligned}$$

By Lemma B.1 and Lemma B.2

$$\begin{aligned} \mathbb{E}_{\nu_t^s} J_t^s(\mathbf{x}, \boldsymbol{\theta}) - \mathbb{E}_{\nu_0^s} J_0^s(\mathbf{x}, \boldsymbol{\theta}_0) &\leq C(\|\nu_t\|_\infty^2, \|\tau_t\|_2^2; \alpha) \mathcal{W}_2(\nu_t^s, \nu_0^s) \\ &\leq C(\|\nu_t\|_\infty^2, \|\tau_t\|_2^2; \alpha) \sqrt{\text{KL}(\nu_t^s \parallel \nu_0^s)}. \end{aligned}$$

For $t \in [0, t_{\max})$, by Definition C.4, we have $\|\nu_t\|_\infty^2, \|\tau_t\|_2^2 = O(d)$. We have

$$\mathbb{E}_{\nu_t^s} J_t^s(\mathbf{x}, \boldsymbol{\theta}) - \mathbb{E}_{\nu_0^s} J_0^s(\mathbf{x}, \boldsymbol{\theta}_0) \leq C(d, \alpha) \sqrt{\text{KL}(\nu_t^s \parallel \nu_0^s)}.$$

Since $\mathbf{p}_{\nu_0}(\mathbf{x}, s) = \mathbf{0}$,

$$\mathbb{E}_{\nu_0^s} J_0^s(\mathbf{x}, \boldsymbol{\theta}) = \mathbf{0}.$$

Therefore,

$$\begin{aligned} \frac{\partial \text{KL}(\nu_t^s \parallel \nu_0^s)}{\partial t} &= \beta \cdot \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} (f_{\tau_t, \nu_t}(\mathbf{x}) - y(\mathbf{x})) \mathbb{E}_{\nu_t^s} J_t^s(\mathbf{x}, \boldsymbol{\theta}) \\ &= \beta \cdot \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} (f_{\tau_t, \nu_t}(\mathbf{x}) - y(\mathbf{x})) \mathbb{E}_{\nu_t^s} (J_t^s(\mathbf{x}, \boldsymbol{\theta}) - J_0^s(\mathbf{x}, \boldsymbol{\theta})) \\ &\leq \beta \cdot C(d, \alpha) \sqrt{\text{KL}(\nu_t^s \parallel \nu_0^s)} \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} (f_{\tau_t, \nu_t}(\mathbf{x}) - y(\mathbf{x})) \\ &\leq \beta \cdot C(d, \alpha) \sqrt{\text{KL}(\nu_t^s \parallel \nu_0^s)} \sqrt{\mathbb{E}(f_{\tau_t, \nu_t}(\mathbf{x}) - y(\mathbf{x}))^2} \\ &= \beta \cdot C(d, \alpha) \sqrt{\text{KL}(\nu_t^s \parallel \nu_0^s)} \sqrt{\widehat{L}(\tau_t, \nu_t)}, \end{aligned}$$

where the last inequality holds owing to the Jensen's inequality.

By the relation $d2\sqrt{x} = dx/\sqrt{x}$,

$$d \left(2\sqrt{\text{KL}(\nu_t^s \parallel \nu_0^s)} \right) \leq \beta \cdot C(d, \alpha) \sqrt{L(\nu_t, \tau_t)} dt.$$

We have for $t \in [0, t_{\max})$,

$$\widehat{L}(\tau_t, \nu_t) \leq e^{-\frac{\beta^2 \Lambda}{2n} t} \widehat{L}(\tau_0, \nu_0).$$

Hence,

$$2\sqrt{\text{KL}(\nu_t^s \|\nu_0^s)} \leq \beta \cdot C(d, \alpha) \int_0^t \sqrt{\widehat{L}(\nu_{t_0}, \tau_{t_0})} dt_0 \leq \frac{4C(d, \alpha)}{\Lambda\beta} \sqrt{\widehat{L}(\tau_0, \nu_0)}.$$

Since $\tau_0(\mathbf{u})$ is standard normal distribution, we have

$$f_{\tau_0, \nu_0}(\mathbf{x}) = \beta \cdot \mathbf{a}^\top \int_{\mathbb{R}^{k_\tau} \times \mathbb{R}^{k_\tau} \times \mathbb{R}} \mathbf{u}^\top \boldsymbol{\sigma}_0(\mathbf{w}^\top \mathbf{Z}_{\nu_0}(\mathbf{x}, 1) + b) d\tau_0(\mathbf{u}, \mathbf{w}, b) = 0,$$

and $|y(\mathbf{x})| \leq 1$ (Assumption 3.1), we have $\widehat{L}(\tau_0, \nu_0) \leq 1$. Therefore, we obtain

$$\text{KL}(\nu_t^s \|\nu_0^s) \leq \frac{C_{\text{KL}}(d, \alpha)}{\Lambda^2 \beta^2},$$

where C_{KL} is a constant dependent only on d, α . \square

Lemma C.8. *Assume the PDE (12) has solution $\tau_t \in \mathcal{P}^2$, and the PDE (10) has solution $\tau_t \in \mathcal{P}^2$. Under Assumption 3.3, 3.1, 3.2, then for all $t \in [0, t_{\max})$, the following results hold:*

$$\text{KL}(\tau_t \|\tau_0) \leq \frac{1}{\Lambda^2 \beta^2} C_{\text{KL}}(d, \alpha).$$

Proof of Lemma C.8. By Gaussian initialization of τ_0 , $\log \tau_0(\boldsymbol{\omega}) = -\frac{\|\boldsymbol{\omega}\|_2^2}{2} + C$, and we have $\nabla_{\boldsymbol{\omega}} \frac{\partial \text{KL}(\tau_t \|\tau_0)}{\partial \tau_t} = \nabla_{\boldsymbol{\omega}} \log \tau_t + \boldsymbol{\omega}$. Therefore, by Lemma 4.6 and Eq. (14), we have

$$\begin{aligned} & \frac{\partial \text{KL}(\tau_t \|\tau_0)}{\partial t} \\ &= -\beta \cdot \int_{\mathbb{R}^{k_\tau}} (\mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n}(f_{\tau_t, \nu_t}(\mathbf{x}) - y(\mathbf{x})) \nabla_{\boldsymbol{\omega}} h(\mathbf{Z}_{\nu_t}(\mathbf{x}, 1), \boldsymbol{\omega})) \cdot (\nabla_{\boldsymbol{\omega}} \log \tau_t + \boldsymbol{\omega}) d\nu_t^s(\boldsymbol{\theta}). \end{aligned}$$

Define

$$\tilde{\mathbf{u}}_t(\boldsymbol{\theta}) := \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} [(f_{\tau_t, \nu_t}(\mathbf{x}) - y(\mathbf{x})) \nabla_{\boldsymbol{\omega}} h(\mathbf{Z}_{\nu_t}(\mathbf{x}, 1), \boldsymbol{\omega})],$$

we have

$$\frac{\partial \text{KL}(\tau_t \|\tau_0)}{\partial t} = -\alpha \cdot \int_{\mathbb{R}^{k_\tau}} \tau_t \tilde{\mathbf{u}}_t \cdot (\nabla_{\boldsymbol{\omega}} \log \tau_t + \boldsymbol{\omega}) d\boldsymbol{\omega} = -\alpha \cdot \int_{\mathbb{R}^{k_\tau}} \tau_t [\tilde{\mathbf{u}}_t \cdot \boldsymbol{\omega} - \nabla_{\boldsymbol{\omega}} \cdot \tilde{v}_t^s] d\boldsymbol{\omega}.$$

We also define

$$I_t(\mathbf{x}, \boldsymbol{\omega}) := -\left(\nabla_{\boldsymbol{\omega}} h(\mathbf{Z}_{\nu_t}(\mathbf{x}, 1), \boldsymbol{\omega}) \cdot \boldsymbol{\omega} - \Delta_{\boldsymbol{\omega}} h(\mathbf{Z}_{\nu_t}(\mathbf{x}, 1), \boldsymbol{\omega}) \right).$$

By the definition of $I_t(\mathbf{x}, \boldsymbol{\omega})$,

$$\frac{\partial \text{KL}(\tau_t \|\tau_0)}{\partial t} = \beta \cdot \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} [(f_{\tau_t, \nu_t}(\mathbf{x}) - y(\mathbf{x})) \mathbb{E}_{\boldsymbol{\omega} \sim \tau_t} I_t(\mathbf{x}, \boldsymbol{\omega})].$$

Similar to the estimate J_t^s , we have the estimation of I_t as

$$\mathbb{E}_{\tau_t} I_t(\mathbf{x}, \boldsymbol{\omega}) - \mathbb{E}_{\tau_0} I_0(\mathbf{x}, \boldsymbol{\omega}) \leq C(\|\tau_t\|_2^2, \|\nu_t\|_\infty^2; \alpha) \sqrt{\text{KL}(\tau_t \|\tau_0)}.$$

and we can have, by setting $\boldsymbol{\omega} = (a, \mathbf{w}, b)$, we have

$$\begin{aligned} \mathbb{E}_{\tau_0} I_0(\mathbf{x}, \boldsymbol{\omega}) &= \mathbb{E}_{\tau_0} (-\nabla_{\boldsymbol{\omega}} h(\mathbf{x}, \boldsymbol{\omega}) \cdot \boldsymbol{\omega} + \Delta_{\boldsymbol{\omega}} h(\mathbf{x}, \boldsymbol{\omega})) \\ &= \mathbb{E}_{(a, \mathbf{w}, b) \sim \mathcal{N}(0, I)} (-a\sigma_0(\mathbf{w}^\top \mathbf{x} + b) - a(\nabla_{\mathbf{w}} \sigma_0(\mathbf{w}^\top \mathbf{x} + b))\mathbf{w} - ab\sigma_0'(\mathbf{w}^\top \mathbf{x} + b)) \\ &\quad + \mathbb{E}_{(a, \mathbf{w}, b) \sim \mathcal{N}(0, I)} (\Delta_{\mathbf{w}} \sigma_0(\mathbf{w}^\top \mathbf{x} + b) + a\sigma_0''(\mathbf{w}^\top \mathbf{x} + b)) = \mathbf{0}. \end{aligned}$$

Therefore, we obtain the KL divergence in a similar fashion.

$$\text{KL}(\tau_t \|\tau_0) \leq \frac{C_{\text{KL}}(d, \alpha)}{\Lambda^2 \beta^2}.$$

\square

Lemma C.9 (Lower bound on the KL divergence). *For any $\tau, \tau' \in \mathcal{P}^2$, $\nu, \nu' \in \mathcal{C}(\mathcal{P}^2; [0, 1])$, if τ', ν' satisfy the Talagrand inequality $T(\frac{1}{2})$, (ref Lemma B.2) We have the lower bound for the KL divergence of τ, τ' and ν, ν' , such that for constant $C_{\text{low}}(\|\tau\|_2^2, \|\tau'\|_2^2, \|\nu\|_\infty^2, \|\nu'\|_\infty^2; \alpha)$,*

$$\sqrt{\text{KL}(\tau\|\tau')} + \sqrt{\text{KL}(\nu\|\nu')} \geq \frac{\mathbb{E}_{\omega \sim \tau} h(\mathbf{Z}_\nu(\mathbf{x}, 1), \omega) - \mathbb{E}_{\omega' \sim \tau'} h(\mathbf{Z}_{\nu'}(\mathbf{x}, 1), \omega')}{C_{\text{low}}(\|\tau\|_2^2, \|\tau'\|_2^2, \|\nu\|_\infty^2, \|\nu'\|_\infty^2; \alpha)}.$$

Proof of Lemma C.9. We have the following estimation

$$\begin{aligned} & \mathbb{E}_{\omega \sim \tau} h(\mathbf{Z}_\nu(\mathbf{x}, 1), \omega) - \mathbb{E}_{\omega' \sim \tau'} h(\mathbf{Z}_{\nu'}(\mathbf{x}, 1), \omega') \\ &= \underbrace{(\mathbb{E}_{\omega \sim \tau} h(\mathbf{Z}_\nu(\mathbf{x}, 1), \omega) - \mathbb{E}_{\omega' \sim \tau'} h(\mathbf{Z}_\nu(\mathbf{x}, 1), \omega'))}_{(A)} \\ & \quad + \underbrace{(\mathbb{E}_{\omega' \sim \tau'} h(\mathbf{Z}_\nu(\mathbf{x}, 1), \omega) - \mathbb{E}_{\omega' \sim \tau'} h(\mathbf{Z}_{\nu'}(\mathbf{x}, 1), \omega'))}_{(B)}, \end{aligned}$$

By Lemma B.4, we have

$$\|\nabla_{\omega} h(\mathbf{Z}_\nu(\mathbf{x}, 1), \omega)\|_2 \leq C_{\sigma} \cdot (\|\mathbf{Z}_\nu(\mathbf{x}, 1)\|_2 + 1) \cdot (\|\omega\|_2 + 1),$$

and by Lemma B.1, Lemma B.6 and Lemma B.2, we have

$$\begin{aligned} (A) &\leq C_{\sigma} \cdot (\|\mathbf{Z}_\nu(\mathbf{x}, 1)\|_2 + 1) \cdot \max\{\|\tau\|_2^2, \|\tau'\|_2^2\} \cdot \mathcal{W}_2(\tau, \tau') \\ &\leq C_{\sigma} \cdot (C_{\mathbf{Z}}(\|\nu_1\|_\infty^2; \alpha) + 1) \cdot \max\{\|\tau\|_2^2, \|\tau'\|_2^2\} \mathcal{W}_2(\tau, \tau'). \\ &\leq 2C_{\sigma} \cdot (C_{\mathbf{Z}}(\|\nu_1\|_\infty^2; \alpha) + 1) \cdot \max\{\|\tau\|_2^2, \|\tau'\|_2^2\} \sqrt{\text{KL}(\tau\|\tau')}. \end{aligned}$$

Besides, by Lemma B.5 and Lemma B.2, we have

$$\begin{aligned} (B) &\leq \mathbb{E}_{\omega' \sim \tau'} C_{\sigma} (\|\omega\|_2^2 + 1) \cdot (\|\mathbf{Z}_\nu(\mathbf{x}, 1) - \mathbf{Z}_{\nu'}(\mathbf{x}, 1)\|_2) \\ &\leq (\|\tau'\|_2^2 + 1) \cdot C_{\mathbf{Z}}(\|\nu_1\|_\infty^2, \|\nu_2\|_\infty^2; \alpha) \cdot \mathcal{W}_2(\nu_1, \nu_2) \\ &\leq 2(\|\tau'\|_2^2 + 1) \cdot C_{\mathbf{Z}}(\|\nu_1\|_\infty^2, \|\nu_2\|_\infty^2; \alpha) \cdot \sqrt{\text{KL}(\nu\|\nu')}. \end{aligned}$$

We let C_{low} be

$$\begin{aligned} & C_{\text{low}}(\|\tau\|_2^2, \|\tau'\|_2^2, \|\nu\|_\infty^2, \|\nu'\|_\infty^2; \alpha) \\ &= \min\{2C_{\sigma} \cdot (C_{\mathbf{Z}}(\|\nu_1\|_\infty^2; \alpha) + 1) \cdot \max\{\|\tau\|_2^2, \|\tau'\|_2^2\}, 2(\|\tau'\|_2^2 + 1) \cdot C_{\mathbf{Z}}(\|\nu_1\|_\infty^2, \|\nu_2\|_\infty^2; \alpha)\} \end{aligned}$$

Therefore, we have

$$\sqrt{\text{KL}(\tau\|\tau')} + \sqrt{\text{KL}(\nu\|\nu')} \geq \frac{\mathbb{E}_{\omega \sim \tau} h(\mathbf{Z}_\nu(\mathbf{x}, 1), \omega) - \mathbb{E}_{\omega' \sim \tau'} h(\mathbf{Z}_{\nu'}(\mathbf{x}, 1), \omega')}{C_{\text{low}}(\|\tau\|_2^2, \|\tau'\|_2^2, \|\nu\|_\infty^2, \|\nu'\|_\infty^2; \alpha)}$$

Since $\|\tau\|_2^2, \|\tau'\|_2^2, \|\nu\|_\infty^2, \|\nu'\|_\infty^2 = O(d)$, we have that the average movement of the KL divergence is on the same order as the change in output value. \square

Lemma C.10. *Assume the PDE (12) has solution $\tau_t \in \mathcal{P}^2$, and the PDE (10) has solution $\nu_t \in \mathcal{P}^2$. Under Assumption 3.3, 3.1, 3.2, then for all $t \in [0, t_{\max})$, We have the lower bound for the KL divergence of τ_t and ν_t , we have for constant $C_{\text{low}}(d; \alpha)$, such that*

$$\sqrt{\text{KL}(\tau_t\|\tau_0)} + \sqrt{\text{KL}(\nu_t\|\nu_0)} \geq \frac{\mathbb{E}_{\omega \sim \tau_t} h(\mathbf{Z}_{\nu_t}(\mathbf{x}, 1), \omega)}{C_{\text{low}}(d; \alpha)}.$$

Proof of Lemma C.10. By the definition of $r_{\max} \leq \sqrt{d}$, and the proof of Lemma C.2, we have $\|\tau_t\|_2^2, \|\nu_t\|_\infty^2 = O(d)$, and we can directly obtain $\|\tau_0\|_2^2 = d + 2, \|\nu_0\|_\infty^2 = 2d + 1$, and $\mathbb{E}_{\omega_0 \sim \tau_0} h(\mathbf{Z}_{\nu_0}(\mathbf{x}, 1), \omega_0) = 0$. Besides, Gaussian initialization satisfies $T(\frac{1}{2})$ condition in Lemma B.2. We have, by Lemma C.9,

$$\sqrt{\text{KL}(\tau_t\|\tau_0)} + \sqrt{\text{KL}(\nu_t\|\nu_0)} \geq \frac{\mathbb{E}_{\omega \sim \tau_t} h(\mathbf{Z}_{\nu_t}(\mathbf{x}, 1), \omega)}{C_{\text{low}}(d; \alpha)},$$

where $C_{\text{low}}(d; \alpha)$ is a constant depending on d, α derived from $C_{\text{low}}(\|\tau\|_2^2, \|\tau'\|_2^2, \|\nu\|_\infty^2, \|\nu'\|_\infty^2; \alpha)$. \square

Lemma C.11. *Under the Assumptions in Lemma C.7, let $\bar{\beta} \geq \frac{4\sqrt{C_{\text{KL}}(d, \alpha)}}{\Lambda r_{\text{max}}}$, we have $t_{\text{max}} = \infty$.*

Proof of Lemma C.11. Otherwise, we have the following inequality, for $\forall t < t_{\text{max}}$:

$$W_2(\nu_t^s, \nu_0^s) \leq 2\sqrt{\text{KL}(\nu_t^s \|\nu_0^s)} \leq \frac{2}{\Lambda \bar{\beta}} \sqrt{C_{\text{KL}}(d, \alpha)}, \quad \forall s \in [0, 1].$$

Therefore,

$$W_2(\nu_t, \nu_0) \leq \frac{2}{\Lambda \bar{\beta}} \sqrt{C_{\text{KL}}(d, \alpha)}.$$

According to the definition of t_{max} , we have $W_2(\nu_t, \nu_0) \leq r_{\text{max}}$. Let

$$\bar{\beta} \geq \frac{4\sqrt{C_{\text{KL}}(d, \alpha)}}{\Lambda r_{\text{max}}}$$

we have $W_2(\nu_t, \nu_0) \leq r_{\text{max}}/2, \forall t \in [0, t_{\text{max}})$, which contradict to the definition of t_{max} in Definition C.4. \square

Proof of Theorem 4.7. Combine the results of Lemma C.11, Lemma C.8, and Lemma C.7, we prove the theorem. \square

C.5 RADEMACHER COMPLEXITY

Proof of Lemma 4.8. Let γ be a parameter whose value will be determined later in the proof. Let $\eta_i, 1 \leq i \leq n$ be the i.i.d. Rademacher random variables,

$$\begin{aligned} \mathcal{R}_n(\mathcal{F}_{\text{KL}}(r)) &= \frac{\beta}{\gamma} \cdot \mathbb{E}_\eta \left(\sup_{\tau: \text{KL}(\tau \|\tau_0) \leq r, \nu: \text{KL}(\nu \|\nu_0) \leq r} \mathbb{E}_\tau \left(\frac{\gamma}{n} \sum_{i=1}^n \eta_i h(\mathbf{Z}_\nu(\mathbf{x}_i, 1), \boldsymbol{\omega}) \right) \right) \\ &\leq \frac{\beta}{\gamma} \cdot \left(r + \mathbb{E}_\eta \sup_{\nu: \text{KL}(\nu \|\nu_0) \leq r} \log \mathbb{E}_{\tau_0} \exp \left(\frac{\gamma}{n} \sum_{i=1}^n \eta_i h(\mathbf{Z}_\nu(\mathbf{x}_i, 1), \boldsymbol{\omega}) \right) \right) \\ &\leq \frac{\beta}{\gamma} \cdot \left(r + \mathbb{E}_\eta \log \mathbb{E}_{\tau_0} \exp \left(\frac{\gamma}{n} \sup_{\nu: \text{KL}(\nu \|\nu_0) \leq r} \sum_{i=1}^n \eta_i h(\mathbf{Z}_\nu(\mathbf{x}_i, 1), \boldsymbol{\omega}) \right) \right) \\ &\leq \frac{\beta}{\gamma} \cdot \left(r + \log \mathbb{E}_{\tau_0} \mathbb{E}_\eta \exp \left(\frac{\gamma}{n} \sup_{\nu: \text{KL}(\nu \|\nu_0) \leq r} \sum_{i=1}^n \eta_i h(\mathbf{Z}_\nu(\mathbf{x}_i, 1), \boldsymbol{\omega}) \right) \right), \end{aligned}$$

where the first inequality is followed by the Donsker-Varadhan representation of KL-divergence in Lemma B.3. The second inequality follows from the increasing function $\log(\cdot), \exp(\cdot)$, and $\sup_y \mathbb{E}_x f(x, y) \leq \mathbb{E}_x \sup_y f(x, y)$, for general variable x, y and function f . The third inequality follows from $\log(\cdot)$'s convexity.

By Assumption 3.3, where $\boldsymbol{\omega} = (a, \mathbf{w}, b)$, we have

$$|h(z_1, \boldsymbol{\omega}) - h(z_2, \boldsymbol{\omega})| \leq C_1 \cdot \|z_1 - z_2\|_2 \cdot a \|\mathbf{w}\|_2.$$

We further estimate

$$\begin{aligned} &\left| \sum_{i=1}^n \eta_i h(\mathbf{Z}_\nu(\mathbf{x}_i, 1), \boldsymbol{\omega}) - \sum_{i=1}^n \eta_i h(\mathbf{Z}_{\nu_0}(\mathbf{x}_i, 1), \boldsymbol{\omega}) \right| \\ &\leq C_1 n \cdot a \|\mathbf{w}\|_2 \cdot \|\mathbf{Z}_\nu(\mathbf{x}_i, 1) - \mathbf{Z}_{\nu_0}(\mathbf{x}_i, 1)\|_2 \\ &\leq C_1 n \cdot a \|\mathbf{w}\|_2 \cdot C_{\mathbf{Z}} (\|\nu\|_\infty^2, \|\nu_0\|_\infty^2; \alpha) \cdot W_2(\nu, \nu_0) \\ &\leq C_1 n \cdot a \|\mathbf{w}\|_2 \cdot C_{\mathbf{Z}} (\|\nu\|_\infty^2, \|\nu_0\|_\infty^2; \alpha) \cdot 2\sqrt{r} \end{aligned}$$

given $\text{KL}(\nu \|\nu_0) \leq r$. Further, we have $\|\nu_0\|_\infty^2 = 2d + 1$, and we have $\|\nu\|_\infty^2 \leq 6d + 2$.

In the following, we use $C_d := 2C_1 \cdot C_Z(6d + 2, 2d + 1; \alpha)$ to denote the constant.

$$\begin{aligned}
\mathcal{R}_n(\mathcal{F}_{\text{KL}}(r)) &\leq \frac{\beta}{\gamma} \cdot \left(r + \log \mathbb{E}_{\tau_0} \mathbb{E}_{\eta} \exp \left(\frac{\gamma}{n} \sum_{i=1}^n \eta_i h(\mathbf{Z}_{\nu_0}(\mathbf{x}_i, 1), \boldsymbol{\omega}) + \gamma \cdot C_d \sqrt{r} \cdot a \|\mathbf{w}\|_2 \right) \right) \\
&= \frac{\beta}{\gamma} \cdot \left(r + \log \mathbb{E}_{\tau_0} \exp \left(\gamma \cdot C_d \sqrt{r} \cdot a \|\mathbf{w}\|_2 \right) \mathbb{E}_{\eta} \exp \left(\frac{\gamma}{n} \sum_{i=1}^n \eta_i h(\mathbf{Z}_{\nu_0}(\mathbf{x}_i, 1), \boldsymbol{\omega}) \right) \right) \\
&\leq \frac{\beta}{\gamma} \cdot \left(r + \log \mathbb{E}_{\tau_0} \exp \left(\gamma \cdot C_d \sqrt{r} \cdot a \|\mathbf{w}\|_2 + \frac{\gamma^2}{2n^2} \sum_{i=1}^n h^2(\mathbf{Z}_{\nu_0}(\mathbf{x}_i, 1), \boldsymbol{\omega}) \right) \right) \\
&\leq \frac{\beta}{\gamma} \cdot \left(r + \frac{1}{2} \log \mathbb{E}_{\tau_0} \exp \left(2\gamma \cdot C_d \sqrt{r} \cdot a \|\mathbf{w}\|_2 \right) + \frac{1}{2} \log \mathbb{E}_{\tau_0} \exp \left(\frac{\gamma^2}{n^2} \sum_{i=1}^n h^2(\mathbf{Z}_{\nu_0}(\mathbf{x}_i, 1), \boldsymbol{\omega}) \right) \right)
\end{aligned}$$

where the first inequality follows from the previous bound; and the second inequality follows from the tail bound: $\mathbb{E}_{\eta} \exp(\sum_{i=1}^n \alpha_i \eta_i) \leq \exp(\frac{1}{2} \sum_{i=1}^n \alpha_i^2)$; and the last inequality follows from the Cauchy inequality.

Still, we set the decomposition $\boldsymbol{\omega} = (a, \mathbf{w}, b) \in \mathbb{R}^{d+2}$, we have $|h(\mathbf{Z}_{\nu_0}(\mathbf{x}_i, 1), \boldsymbol{\omega})| = |a\sigma_0(\mathbf{w}^\top \mathbf{Z}_{\nu_0}(\mathbf{x}_i, 1) + b)| \leq |a|C_1$. We have

$$\begin{aligned}
&\mathcal{R}(\mathcal{F}_{\text{KL}}(r)) \\
&\leq \frac{\beta}{\gamma} \cdot \left(r + \frac{1}{2} \log \mathbb{E}_{(a, \mathbf{w}, b) \sim \mathcal{N}(0, \mathbf{I})} \exp \left(2\gamma \cdot C_d \sqrt{r} \cdot a \|\mathbf{w}\|_2 \right) + \frac{1}{2} \log \mathbb{E}_{a \sim \mathcal{N}(0, 1)} \exp \left(\frac{\gamma^2 a^2 C_1^2}{n} \right) \right)
\end{aligned}$$

We remark that

$$\begin{aligned}
\log \mathbb{E}_{t \sim \mathcal{N}(0, 1)} \exp(Ct^2) &= -\frac{1}{2} \log(1 - 2C) \leq 2C \\
\log \mathbb{E}_{t \sim \mathcal{N}(0, 1)} \exp(Ct) &= \exp(C^2/2).
\end{aligned}$$

where the first inequality holds for $C \leq 1/4$.

Therefore, setting $\gamma = \sqrt{nr}/C_1$, we have

$$\begin{aligned}
\frac{1}{2} \log \mathbb{E}_{a \sim \mathcal{N}(0, 1)} \exp \left(\frac{\gamma^2 a^2 C_1^2}{n} \right) &= -\frac{1}{4} \log(1 - 2r) \leq r, \\
\frac{1}{2} \log \mathbb{E}_{(a, \mathbf{w}, b) \sim \mathcal{N}(0, \mathbf{I})} \exp \left(2\gamma \cdot C_d \sqrt{r} \cdot a \|\mathbf{w}\|_2 \right) &= \frac{1}{2} \log \mathbb{E}_{\boldsymbol{\omega} \sim \mathcal{N}(0, \mathbf{I})} \exp(2\gamma^2 C_d^2 r \|\mathbf{w}\|_2^2) \\
&\leq -\frac{1}{4} \log(1 - 4nr^2(C_d/C_1)^2) \leq 2nr^2(C_d/C_1)^2.
\end{aligned}$$

where the inequality holds iff $r \leq \frac{1}{4}$ and $nr^2(C_d/C_1)^2 \leq \frac{1}{8}$. We set $r_0 = \min\{1/4, 1/(4\sqrt{n}) \cdot C_1/C_d\}$, we have for $\forall r \leq r_0$,

$$\mathcal{R}(\mathcal{F}_{\text{KL}}(r)) \leq \beta/\gamma(r + 2nr^2(C_d/C_1)^2 + r) \leq \beta \cdot \sqrt{r/n} \cdot 2(C_1 + C_d).$$

where \lesssim hides constant. \square

Theorem C.12 (Rademacher complexity). For any $\delta > 0$, with probability at least $1 - \delta$, the following bound holds $\forall f_{\tau, \nu} \in \mathcal{F}_{\text{KL}}(r)$:

$$\mathbb{E}_{\mathbf{x} \sim \mu_{\mathbf{x}}} \ell_{0-1}(f_{\tau, \nu}(\mathbf{x}), y(\mathbf{x})) \leq 4\mathcal{R}_n(\mathcal{F}_{\text{KL}}(r)) + 6\sqrt{\log(2/\delta)/2n} + \sqrt{\mathbb{E}_{\mathcal{D}_n}(f_{\tau, \nu}(\mathbf{x}) - y(\mathbf{x}))^2}$$

Proof of Theorem C.12. We introduce the additional loss function

$$\bar{\ell}(f, y) = \max\{\min\{1 - 2yf, 1\}, 0\}.$$

By definition, we have $\bar{\ell}$ is 2-Lipschitz in the first argument, and

$$\ell_{0-1}(f, y) \leq \bar{\ell}(f, y) \leq |f - y|,$$

for any $f \in \mathbb{R}$ and $y \in \{\pm 1\}$. Using the standard properties of Rademacher complexity, we have that with probability at least $1 - \delta$, for all $f \in \mathcal{F}_{\text{KL}}(r)$

$$\mathbb{E}_{\mu_X} \bar{\ell}(f, y) \leq \mathbb{E}_{\mathcal{D}_n} \bar{\ell}(f, y) + 4\mathcal{R}_n(\mathcal{F}_{\text{KL}}(r)) + 6\sqrt{\frac{\log(2/\delta)}{2n}}.$$

Therefore, we have

$$\mathbb{E}_{\mu_X} \ell_{0-1}(f, y) \leq \mathbb{E}_{\mu_X} \bar{\ell}(f, y) \leq \sqrt{\mathbb{E}_{\mathcal{D}_n} (f - y)^2} + 4\mathcal{R}_n(\mathcal{F}_{\text{KL}}(r)) + 6\sqrt{\frac{\log(2/\delta)}{2n}}$$

□

Lemma C.13. Let $\tau_y \in \mathcal{C}(\mathcal{P}^2; [0, 1])$ and $\nu_y \in \mathcal{P}^2$ be the ground truth distributions, such that,

$$y(\mathbf{x}) := \mathbb{E}_{\boldsymbol{\omega} \sim \tau_y} h(\mathbf{Z}_{\nu_y}(\mathbf{x}, 1), \boldsymbol{\omega}).$$

Then, we have the bound for the KL divergence, for τ_*, ν_* satisfying $\widehat{L}(\tau_*, \nu_*) = 0$,

$$\max\{\text{KL}(\tau_* \|\tau_0), \text{KL}(\nu_* \|\nu_0)\} \leq \beta^{-2}(\chi^2(\tau_y \|\tau_0) + \chi^2(\nu_y \|\nu_0)).$$

Proof of Lemma C.13. We assume that $\{\tau_*^\lambda, \nu_*^\lambda\}$ is the solution to the following minimization problem:

$$\{\tau_*^\lambda, \nu_*^\lambda\} = \arg \min_{\tau, \nu} \widehat{L}(\tau, \nu) + \lambda(\text{KL}(\tau \|\tau_0) + \text{KL}(\nu \|\nu_0)).$$

Consider the mixture distribution $\widehat{\tau}, \widehat{\nu}$ be defined as

$$(\widehat{\tau}, \widehat{\nu}) = \frac{\beta - 1}{\beta}(\tau_0, \nu_0) + \frac{1}{\beta}(\tau_y, \nu_y),$$

and we have

$$\begin{aligned} \widehat{L}(\widehat{\tau}, \widehat{\nu}) &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} \left(\frac{\beta - 1}{\beta} \cdot \beta \mathbb{E}_{\boldsymbol{\omega} \sim \tau_0} h(\mathbf{Z}_{\nu_0}(\mathbf{x}, 1), \boldsymbol{\omega}) + \frac{1}{\beta} \cdot \beta \mathbb{E}_{\boldsymbol{\omega} \sim \tau_y} h(\mathbf{Z}_{\nu_y}(\mathbf{x}, 1), \boldsymbol{\omega}) - y(\mathbf{x}) \right) \\ &= \mathbb{E}_{\mathbf{x} \sim \mathcal{D}_n} \left(0 + \frac{1}{\beta} \cdot \beta \cdot y - y \right)^2 = 0, \end{aligned}$$

and by the definition of $\tau_*^\lambda, \nu_*^\lambda$, we obtain

$$\widehat{L}(\tau_*^\lambda, \nu_*^\lambda) + \lambda(\text{KL}(\tau_*^\lambda \|\tau_0) + \text{KL}(\nu_*^\lambda \|\nu_0)) \leq \widehat{L}(\widehat{\tau}, \widehat{\nu}) + \lambda(\text{KL}(\widehat{\tau} \|\tau_0) + \text{KL}(\widehat{\nu} \|\nu_0)).$$

This leads to

$$\begin{aligned} \widehat{L}(\tau_*^\lambda, \nu_*^\lambda) &\leq \lambda(\text{KL}(\widehat{\tau} \|\tau_0) + \text{KL}(\widehat{\nu} \|\nu_0)) \\ \text{KL}(\tau_*^\lambda \|\tau_0) + \text{KL}(\nu_*^\lambda \|\nu_0) &\leq \text{KL}(\widehat{\tau} \|\tau_0) + \text{KL}(\widehat{\nu} \|\nu_0). \end{aligned}$$

Taking $\lambda \rightarrow 0$, we have $\tau_*^\lambda \rightarrow \tau_*$, and $\nu_*^\lambda \rightarrow \nu_*$. Accordingly, we have

$$\begin{aligned} \widehat{L}(\tau_*, \nu_*) &= 0 \\ \text{KL}(\tau_* \|\tau_0) + \text{KL}(\nu_* \|\nu_0) &\leq \text{KL}(\widehat{\tau} \|\tau_0) + \text{KL}(\widehat{\nu} \|\nu_0) \end{aligned}$$

We can explicitly compute the KL divergence such that

$$\text{KL}(\widehat{\tau} \|\tau_0) \leq \chi^2(\widehat{\tau} \|\tau_0) = \int \left(\frac{\beta - 1}{\beta} + \frac{\tau_y(\boldsymbol{\omega})}{\beta \tau_0(\boldsymbol{\omega})} - 1 \right)^2 d\boldsymbol{\omega} = \beta^{-2} \chi^2(\tau_y \|\tau_0),$$

and similarly, we have

$$\text{KL}(\widehat{\nu} \|\nu_0) \leq \beta^{-2} \chi^2(\nu_y \|\nu_0).$$

Finally, we conclude the proof. □

Proof of Theorem 4.9. By Theorem C.12, for $r > 0$, for any $\delta > 0$, with probability at least $1 - \delta$, the following bound holds $\forall f_{\tau, \nu} \in \mathcal{F}_{\text{KL}}(r)$:

$$\mathbb{E}_{\mathbf{x} \sim \mu_X} \ell_{0-1}(f_{\tau, \nu}(\mathbf{x}), y(\mathbf{x})) \leq 4\mathcal{R}_n(\mathcal{F}_{\text{KL}}(r)) + 6\sqrt{\log(2/\delta)/2n} + \sqrt{\mathbb{E}_{\mathcal{D}_n}(f_{\tau, \nu}(\mathbf{x}) - y(\mathbf{x}))^2}.$$

By Lemma C.13, and the definition of r_0 in Lemma 4.8, we set β such that

$$\beta^{-2}(\chi^2(\tau_y \parallel \tau_0) + \chi^2(\nu_y \parallel \nu_0)) \leq r_0,$$

i.e.,

$$\beta \geq \sqrt{\frac{\chi^2(\tau_y \parallel \tau_0) + \chi^2(\nu_y \parallel \nu_0)}{r_0}} = O(\sqrt{n}),$$

and

$$f_{\tau_*, \nu_*} \in \mathcal{F}_{\text{KL}}(\beta^{-2}(\chi^2(\tau_y \parallel \tau_0) + \chi^2(\nu_y \parallel \nu_0))).$$

Therefore, we apply the Rademacher complexity bound in Lemma 4.8, and we have

$$\mathcal{R}_n(\mathcal{F}_{\text{KL}}(\beta^{-2}(\chi^2(\tau_y \parallel \tau_0) + \chi^2(\nu_y \parallel \nu_0)))) \lesssim \beta \sqrt{\frac{\beta^{-2}(\chi^2(\tau_y \parallel \tau_0) + \chi^2(\nu_y \parallel \nu_0))}{n}} = O(1/\sqrt{n}).$$

where β cancels out. By Lemma C.13, we have $\widehat{L}(\tau_*, \nu_*) = 0$. Finally,

$$\mathbb{E}_{\mathbf{x} \sim \mu_X} \ell_{0-1}(f_{\tau, \nu}(\mathbf{x}), y(\mathbf{x})) \lesssim O(1/\sqrt{n}) + 6\sqrt{\log(2/\delta)/2n}.$$

□

C.6 EXPERIMENTS

We validate our findings on the toy dataset ‘‘Two Spirals’’, where the data dimension $d = 2$. We use a neural ODE model (Poli et al., 2021) to approximate the infinite depth ResNets, where we take the discretization $L = 10$. The neural ODE model and the output layer are both parametrized by a two-layer network with the tanh activation function, and the hidden dimension is $M = K = 20$. The parameters of the ResNet encoder and the output layer are jointly trained by Adam optimizer with an initial learning rate 0.01. We perform full-batch training for 1,000 steps on the training dataset of size n_{train} , and test the resulting model on the test dataset of size $n_{\text{test}} = 1024$ by the 0-1 classification loss. We run experiments over 3 seeds and report the mean. We fit the results (after logarithm) by ordinary least squares, and obtain the slope is -1.02 with p-value 10^{-5} , as shown in Figure 1. That means, the obtained rate is $O(1/n)$, which is faster than our derived $O(1/\sqrt{n})$ rate. We hope we could use some localized schemes, e.g., local Rademacher complexity, to close the gap.




Figure 1: **Left:** ‘‘Two Spirals’’ datasets. **Right:** L_{0-1} test error v.s. the training dataset size n_{train} (blue), OLS fitted line (red) which is close to the $O(1/n)$ rate with p-value 10^{-5} .