arXiv:2205.09860v2 [cs.LG] 18 Oct 2022

# Mean-Field Analysis of Two-Layer Neural Networks: Global Optimality with Linear Convergence Rates

Jingwei Zhang[‡ †], Xunpeng Huang [‡], and Jincheng Yu [‡]

[‡]Department of Computer Science and Engineering, HKUST
[†]Huawei Noah's Ark Lab
{jzhangey, xhuangck, jyubh}@connect.ust.hk

**Abstract**

We consider optimizing two-layer neural networks in the mean-field regime where the learning dynamics of network weights can be approximated by the evolution in the space of probability measures over the weight parameters associated with the neurons. The mean-field regime is a theoretically attractive alternative to the NTK (lazy training) regime which is only restricted locally in the so-called neural tangent kernel space around specialized initializations. Several prior works (Chizat and Bach (2018); Mei et al. (2018)) establish the asymptotic global optimality of the mean-field regime, but it is still challenging to obtain a quantitative convergence rate due to the complicated unbounded nonlinearity of the training dynamics. This work establishes the first linear convergence result for vanilla two-layer neural networks trained by continuous-time noisy gradient descent in the mean-field regime. Our result relies on a novel time-depdendent estimate of the logarithmic Sobolev constants for a family of measures determined by the evolving distribution of hidden neurons.

## 1 Introduction

Gradient-based optimization is a fundamental tool in machine learning and has witnessed great empirical success for training neural networks, despite the highly non-convexity landscape of the objective. However, theoretical understanding of nonconvex optimization in neural networks is quite limited. Until recently, there has been much work that explains the success of gradient-based optimization in overparametrized neural networks, that is neural networks with massive hidden units. Under the overparametrization condition, the learning problem can be translated into minimizing a convex functional and hence circumventing the difficulties of analyzing non-convex objectives.

It's worth mentioning that there has been much broader interest in analyzing the convergence of machine learning algorithms by formulating it as the problem of minimizing some (usually convex) functional of a measure, such as variational inference (Liu and Wang (2016); Liu (2017); Chewi et al. (2020)), generative adversatial networks (Johnson and Zhang (2019); Nitanda and Suzuki (2020)) and learning infinite-width neural networks (Chizat and Bach (2018); Mei et al. (2018); Nguyen and Pham (2020); Fang et al. (2021)). The key idea is by approximating the learning dynamics of model parameters by the optimization on the space of probability of measures over the model parameters under the overparametrization condition. Such functional optimization

procedure corresponds to the gradient flow over the space of probability measure endowed with some probability metrics such as the Wasserstein distance (Chizat and Bach (2018); Mei et al. (2018)). The functional optimization framework enables us to study the optimization and generalization of learning algorithms with a different set of tools, such as non-linear PDEs for optimial transport (Ambrosio et al. (2008); Villani (2009); Santambrogio (2015)).

### 1.1 Related Works

The study of overparametrized neural networks mainly falls into two categories: the NTK regime or the mean-field regime.

**NTK Analysis of Neural Networks.** It has been already known in the literature that deep neural network in the infinite-width limit is equivalent to a Gaussian process since decades ago (Neal (1996); Williams (1997); Winther (2001); Neal (2012); Lee et al. (2018)) and thus they can be connected to kernel methods. Later, Jacot et al. (2018) shows that the evolution of the trainable parameters in continuous width deep neural nets during training process can be captured by a kernel, which is referred to Neural Tangent Kernel (NTK). Recent work has shown that with a specialized scaling and random initialization, the parameters of continuous width two-layer neural network is restricted in an infinitesimal region around the initialization and can be regarded a linear model with infinite dimensional feature (Du et al. (2018); Li and Liang (2018); Arora et al. (2019); Du et al. (2019); Li et al. (2020)). Since the system becomes linear, the dynamics of GD within this region can be tracked via properties of the associated NTK and the convergence to the global optima with a linear conergence rate can be proved. Later, Allen-Zhu et al. (2018); Zou et al. (2018); Allen-Zhu et al. (2019) extend the NTK analysis to the global convergence of GD for multi-layer neural nets. Despite the huge theoretic progress of NTK regime in analyzing both shallow and deep neural networks, such regime deviates far from practical fully trained neural networks. For NTK approximation, the training is restricted to a small region around the initialization and therefore it suffers from lazy training and cannot explain why neural networks learn discriminative features in practice. An alternative and more promising approach is to consider the mean-field modeling of neural networks.

**Mean-Field Analysis of Neural Networks.** Another strand of work tries to study the infinite-width neural network from a mean field view (Chizat and Bach (2018); Mei et al. (2018); Sirignano and Spiliopoulos (2019); Nguyen and Pham (2020); Fang et al. (2021)). The key idea is to characterize the learning dynamics of SGD (Chizat and Bach (2018)) or noisy SGD (Mei et al. (2018)) as the gradient flow over the space of probability distribution of NN parameters. More specifically, it has been proved in Chizat and Bach (2018) that SGD enjoys asymptotic global optimality under homogeneous activations in the mean-field regime. While independently in Mei et al. (2018), the asymptotic global optimality relies on the mixing effect of noise, hence the analysis mainly focuses on the noisy SGD. In this paper, we analyze the noisy SGD and hence adopt the framework in Mei et al. (2018) that the learning dynamics of noisy SGD can be characterized by a non-linear PDE of Mckean-Vlason type, which corresponds to the Wasserstein gradient flow over the space of distributions of NN parameters. When the time goes to infinity, the noisy SGD convergences asymptotically to the unique global optima of the convex objective functional for two-layer NN with infinite width.

Despite lots of advances in analyzing the dynamics and global convergence of neural networks trained by noisy SGD in the mean-field regime (Mei et al. (2018); Chizat and Bach (2018)), these results are qualitative without known convergence rates. Hence, an important research question

arises:

*Can we obtain a quantitative convergence rate for overparametrized neural networks trained by noisy SGD in the mean-field regime?*

This question remains open even for the simple vanilla two-layer neural networks. In this paper, we address this question and prove the first linear convergence result for standard two-layer neural networks trained by continuous-time noisy gradient descent in the mean-field regime.

Prior to the submission of this work, we are aware that several recent independent works (Chizat (2022); Nitanda et al. (2022)) have been on Arxiv that also tried to prove the linear convergence result for neural networks trained by noisy SGD in the mean-field regime, but as pointed out by the authors (Chizat (2022)), their assumptions of boundedness and smoothness in both works cannot be applied to the vanilla two-layer neural networks (we will discuss details in later sections). Our work differs from them in both assumptions and proof techniques. Specifically,

- In Chizat (2022); Nitanda et al. (2022), it is assumed that each particle is bounded and smooth with respect both weight and position of each neuron. This assumption is unrealistic even for vanilla two-layer neural networks unless we fix the parameter of the second layer and assume the activation is bounded. In such case, it is not a real two-layer neural network.

- As we will see later in this work, the unboundedness and non-smoothness of each particle with respect to its parameter is the main technical hurdle to prove the linear convergence result. Unlike Chizat (2022); Nitanda et al. (2022) which directly make restrictive assumptions on the function of neural networks, we address this technical hurdle by imposing regularizers with particular growth rate and hence it can control the curvature of the whole objective. The regularizer is similar to weight decay and can be easily implemented.

- Our technique could yield a more refined time-dependent estimate of the logarithmic sobolev constant and if it is uniformly upper bounded for all $t \geq 0$, then such an upper bound can be the linear convergence rate. While in Chizat (2022); Nitanda et al. (2022), it is assumed that the logarithmic Sobolev inequality uniformly holds for some set of measures in $\mathcal{P}_2(\mathbb{R}^d)$ and the estimate of Sobolev constant is directly obtained from assumptions which is independent of the time. Such rough constant may not reveal the real convergence rate of the algorithm.

There are other existing works that also show the linear convergence of two-layer neural networks in the mean-field regime. More specifically, Chen et al. (2022) show that under the non-degeneracy condition on certain Gram matrix, gradient descent (GD) can linear converge to 0 training loss efficiently in the mean-field scaling. Our approach is not directly comparable to Chen et al. (2022) since the analysis is an extension of NTK-theory to the feature learning regime that is different from the continuous noisy SGD in the mean-field regime studied in our paper. A more related work is Hu et al. (2019) which obtain the linear convergence of noisy SGD for two-layer NN in the mean-field regime, but the convergence result requires the sufficiently strong regularization and does not exploit the convexity of the functional of neural networks. In some sense, the convergence in Hu et al. (2019) is only because of sufficiently strong regularizations. While in our analysis, the regularization could be arbitrarily small.

### 1.2 Contribution

In this paper, we prove that under some mild regularity conditions on the regularizer, the solution of the noisy SGD on a two-layer neural network converges linearly to the global optima of the regularized objective in the mean-field regime. More specifically,

- We present an analysis to show that logarithmic Sobolev inequality can be directly applied to mean-filed analysis of two-layer neural networks. If a uniform upper bound on the logarithmic Sobolev constants can be obtained for a family of measures determined by the evolving distribution of hidden neurons along the trajectory, then we obtain a linear convergence rate in the mean-field regime.
- We develop new techniques to estimate the upper bound of logarithmic Sobolev constants for measures determined by the evolving distribution of parameters of two-layer neural networks. We show that for some carefully chosen regularizer, an upper bound on the logarithmic Sobolev constants can be obtained and hence we can achieve the linear convergence result for the vanilla two-layer neural networks for the first time. This solves the unboundedness issue raised by Chizat (2022), where the authors claim that there should be incompatibility between noisy gradient descent and the standard architecture of two-layer neural networks. We instead show this should be compatible if we properly chose the regularizer.

Our theory also has practical implications: different regularizers could lead to different estimate of the log-Sobolev constants, hence influencing the convergence rates to the global optima. Designing good regularizers for effective training neural networks with theoretic guarantees is a promising direction for future research.

### 1.3 Notations

Let $\|\cdot\|_2$ denote the Euclidean norm of a vector and $\|\cdot\|_{op}$ denote the operator norm of a matrix. We use $\nabla$, $\nabla\cdot$, $\Delta := \nabla\cdot\nabla$ and $\nabla^2$ to denote the gradient, divergence, Laplacian, and Hessian operator respectively. $\langle\cdot,\cdot\rangle$ denotes the inner product between vectors or functions. For $\mu,\nu\in\mathbb{P}_2(\mathbb{R}^d)$, the 2-Wasserstein distance is defined as

$$W_2(\mu,\nu)^2 := \min_{\gamma\in\Pi(\mu,\nu)}\int\|x-y\|_2^2\gamma(dx,dy) \tag{1.1}$$

where $\Pi(\mu,\nu)$ denotes the coupling between $\mu$ and $\nu$, which means $\Pi(\mu,\nu)$ is a measure on $\mathbb{R}^d\times\mathbb{R}^d$ with marginals $\mu$ and $\nu$ respectively.

## 2 Problem Setup

For an input feature vector $x\in\mathbb{R}^d$, consider a two-layer neural network with $N$ hidden neurons

$$\hat{f}_N(x,\Theta) = \frac{1}{N}\sum_{i=1}^N\tilde{h}(\theta_i,x) = \frac{1}{N}\sum_{i=1}^N u_i h(w_i,x)\ . \tag{2.1}$$

Here, $N$ is the number of hidden neurons; $\Theta=[\theta_1,\dots,\theta_N]\in\mathbb{R}^{(d+1)\times N}$ are parameters of the neural network; each $\theta_i=(u_i,w_i)\in\mathbb{R}^{d+1}$ denotes a weighted neuron, with position $w_i\in\mathbb{R}^d$ and weight

$u_i \in \mathbb{R}$; $h(\cdot,\cdot)$ is the (nonlinear) activation function.

In the mean-field regime, each weighted neuron $\theta_i$ can be regarded as a particle that is drawn i.i.d. from a distribution $\rho(\theta)$. When the number of hidden neurons tends to infinity, we obtain the mean-field limit of the two-layer neural network

$$f(x,\rho) = \int_{\mathbb{R}^{d+1}} \tilde{h}(\theta,x)\rho(\theta)d\theta \ . \tag{2.2}$$

In this paper we consider the classical setting of supervised learning whereby we are given a feature-label pair $(x,y) \in \mathbb{R}^d \times \mathbb{R}$ drawn from an unknown distribution. If we choose the loss function $\phi : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$, we wish to minimize the expected loss over the data distribution. In order to control the model capacity, we also add some regularizers to the parameters of the neural network. Then the objective of the two-layer neural network in the mean-field limit becomes:

$$Q(\rho) = \mathbb{E}_{x,y}\phi\left(\int_{\mathbb{R}^{d+1}} \tilde{h}(\theta,x)\rho(\theta)d\theta, y\right) + \int_{\mathbb{R}^{d+1}} r(\theta)\rho(\theta)d\theta + \lambda\int_{\mathbb{R}^{d+1}} \rho(\theta)\log\rho(\theta)d\theta \tag{2.3}$$

where $\tilde{h}(\theta,x) = uh(w,x)$ with $u \in \mathbb{R}$ and $\theta = [u,w] \in \mathbb{R}^{d+1}$; $\rho(\theta)$ is a probability distribution of neurons; $r(\theta)$ is the regularizer that penalizes the weight norm of neurons.

**Remark 1.** *The above formulation of infinite-width two-layer neural networks can be easily implemented by i.i.d sampling the weight $\theta$ of neurons over the distribution $\rho(\theta)$. For the second term, it corresponds to the weight decay when $r(\theta)$ is the square of the Euclidean norm of $\theta$. For the last entropy-regularized term, it can be implemented by adding Gaussian noise, which corresponds to the noise term in the training of the neural network weights by noisy SGD.*

**Remark 2.** *In Chizat (2022) and Nitanda et al. (2022), it is assumed that $\tilde{h}(\theta,x)$ is bounded and smooth with respect to $\theta$ uniformly for any $x$. This assumption is restrictive and cannot be applied in two-layer neural networks, since we could not control the boundedness of $\tilde{h}(\theta,x) = uh(w,x)$ with respect to $u$. Even when we add another bounded and smooth activation $g$ in the output, and let $\tilde{h}(\theta,x) = g(uh(w,x))$, the smoothness of $\tilde{h}(\theta,x) = g(uh(w,x))$ with respect to $\theta$ still cannot hold. As shown in our proof, the spectral norm of the Hessian $\nabla^2_\theta \tilde{h}(\theta,x)$ will tend to infinity as $|u| \to \infty$.*

## 2.1 The Decreasing Path Determined by the Gradient Flow

In order to minize the regularized objective (2.3), we consider the following particle optimization for neurons

$$d\theta_t = -\nabla_\theta g(\theta_t,\rho_t)dt \ . \tag{2.4}$$

It's straightforward to get the following as the standard result of the Fokker-Planck equation:

$$\frac{\partial\rho_t(\theta)}{\partial t} = \nabla\cdot[\rho_t\nabla g(\theta,\rho_t)]. \tag{2.5}$$

Then, we obtain the following decreasing lemma, which is standard and well-known in the existing literature by simply integral-by-parts.

**Lemma 1** (The Decreasing Lemma). *If we take*

$$g(\theta,\rho)=U(\theta,\rho)+\lambda\log\rho(\theta)\quad \text{and}\quad U(\theta,\rho)=\mathbb{E}_{x,y}\phi_1'(f(\rho,x),y)\tilde{h}(\theta,x)+r(\theta) \tag{2.6}$$

*where* $f(\rho,x)=\int_{\mathbb{R}^{d+1}}\tilde{h}(\theta,x)\rho(\theta)d\theta$ *and* $\phi_1'(\cdot,\cdot)$ *is the derivative of* $\phi$ *with respect to its first argument. We can get*

$$\frac{dQ(\rho_t)}{dt}=-\int_{\mathbb{R}^{d+1}}\|\nabla_\theta g(\theta,\rho_t)\|_2^2\rho_t(\theta)d\theta\ . \tag{2.7}$$

*Proof.* By definition and integral by parts, we obtain

$$\frac{dQ(\rho_t)}{dt}=\left\langle\frac{\delta Q(\rho_t)}{\delta\rho_t},\frac{\partial\rho_t(\theta)}{\partial t}\right\rangle=\langle g(\theta,\rho_t),\nabla\cdot[\rho_t\nabla g(\theta,\rho_t)]\rangle=-\left\langle\|\nabla_\theta g(\theta,\rho_t)\|_2^2,\rho_t\right\rangle\ .$$

□

**Remark 3.** *Notice that in Lemma 1, we have picked*

$$g(\theta,\rho)=\mathbb{E}_{x,y}\phi_1'(f(\rho,x),y)\tilde{h}(\theta,x)+r(\theta)+\lambda\log\rho(\theta)\ ,$$

*which is in fact the linear functional derivative of the objective* $Q(\rho)$*. Then the particle optimization for neurons becomes*

$$d\theta_t=-\nabla_\theta U(\theta_t,\rho_t)dt-\lambda\nabla\log\rho_t(\theta)dt \tag{2.8}$$

*and the corresponding nonlinear Fokker-Planck equation becomes*

$$\frac{\partial\rho_t(\theta)}{\partial t}=\nabla\cdot[\rho_t\nabla U(\theta,\rho_t)]+\lambda\Delta\rho_t(\theta). \tag{2.9}$$

*Notice that the particle update (2.8) is hard to implement due to the term* $-\lambda\nabla\log\rho_t(\theta)dt$*, which can be replaced by adding diffusions, i.e.,*

$$d\theta_t=-\nabla_\theta U(\theta_t,\rho_t)dt+\sqrt{2\lambda}dB_t \tag{2.10}$$

*where* $dB_t$ *is the standard Brownian motion.*

**Remark 4.** *The decreasing lemma shows that the risk will be non-increasing along the gradient flow dynamics. It is quite standard in the literature of Wasserstein gradient flows (e.g., see Section* 4.4.*E of Ambrosio and Savaré (2007)) where one needs the existence of the first-order variations and the convexity along the Wasserstein geodesics for the objective functional , which can be easily verified in our case.*

The well-posedness of the nonlinear Fokker-Planck equation for standard two-layer neural networks has been proved in Mei et al. (2019), where the potentials are not globally bounded and Lipschitz anymore.

**Proposition 1** (Generalization of the Proposition 51 in Mei et al. (2019)). *Let the initialization* $\rho_0\in\mathbb{P}_2(\mathbb{R}^{d+1})$ *such that* $Q(p_0)<\infty$*. Then, the solution* $(\rho_t)_{t\geq 0}\in\mathbb{P}_2(\mathbb{R}^{d+1})$ *of (2.9) exists and is unique under some mild smoothness and growth conditions on the regularizer* $r(\theta)$*.*

### 2.2 Assumptions

Before we introduce our main result, here we provide several assumptions on the loss function $\phi(\cdot,\cdot)$ and the activation function $h(\cdot,\cdot)$ that are critical in our subsequent analysis.

**Assumption 1.** *The loss function $\phi(\cdot,\cdot)$ satisfies the following conditions:*

- *(Convexity.) $\phi(\cdot,\cdot)$ is convex on its first argument.*
- *(Lower Boundedness.) $\phi(\cdot,\cdot)$ is bounded below, i.e., $\exists B_l > 0$ such that $\forall y' \in \mathbb{R}$ and $y \in \mathbb{R}$, $\phi(y', y) \geq B_l$ .*
- *(Lipschitz Gradient.) $\phi(\cdot,\cdot)$ is $L_1$-bounded and $L_2$-Lipschitz continuous gradient on its first argument, i.e., $\exists L_1, L_2 > 0$ such that $\forall y_1, y_2 \in \mathbb{R}$ and $y \in \mathcal{Y}$, $|\phi_1'(y_1, y)| \leq L_1$ and*

$$|\phi_1'(y_1, y) - \phi_1'(y_2, y)| \leq L_2|y_1 - y_2| \ .$$

Note that such assumptions on the loss function are commonly used in existing literatures (Du et al. (2019)). We also need some assumptions on the activation function.

**Assumption 2.** *The feature activation function satisfies the following conditions:*

- *(Differentiability.) $\forall x \in \mathcal{X}$, $h(\cdot, x)$ is second-order differentiable.*
- *(Growth Condition.) $\forall x \in \mathcal{X}$ and $w \in \mathbb{R}^d$, $|h(w,x)| \leq C_1\|w\| + C_2$ , $\|\nabla_1 h(w,x)\| \leq C_3$ , and $\|\nabla_1^2 h(w,x)\|_{op} \leq C_4$ for some constants $C_1, C_2, C_3, C_4 > 0$ .*

The assumptions on the activation function also incorporate many commonly used activations, such as sigmoid and smoothed ReLU.

## 3 Global Optimality with Linear Convergence Rates

We need to introduce some functional inequalities that is essential in the proofs of theorems.

### 3.1 Poincare and Logarithmic Sobolev Inequality

Consider a probability space $(\mathcal{X}, \mathcal{F}, \mu_t)$ where $\mu_t$ is some measure at time $t \geq 0$ and let $\mathcal{L}_t$ be a self-adjoint operator on $L^2(\mu_t)$ with domain $\mathbb{D}_2(\mathcal{L}_t)$ which is a dense subset of $L^2(\mu_t)$ and the corresponding Dirichlet form associated to $\mathcal{L}_t$ is

$$\mathcal{E}_t(f) = -\langle \mathcal{L}_t f, f\rangle_{\mu_t}, \quad f \in \mathbb{D}_2(\mathcal{L}_t) \ . \tag{3.1}$$

In our setting, we consider the symmetric diffusion in the Euclidean space. More specifically, we let $\mathcal{X} = \mathbb{R}^{d+1}$ and let $\mathcal{L}_t = \nabla\cdot\nabla - \nabla U_t \cdot \nabla$ for some $U_t \in W^{1,2}_{loc}(\mathbb{R}^{d+1})$ such that $Z = \int_{\mathbb{R}^{d+1}} e^{U_t(\theta)} d\theta < \infty$ and $d\mu_t(\theta) = Z^{-1} e^{-U_t(\theta)} d\theta$. In this case, the associated Dirichlet form is given by the symmetric property as follows

$$\mathcal{E}_t(f) = \int_{\mathbb{R}^{d+1}} \|\nabla f(\theta)\|_2^2 \, d\mu_t(\theta) \ . \tag{3.2}$$

Throughout the paper, if it is not mentioned explicitly, all test functions in the functional inequalities should be in $\mathbb{D}_2(\mathcal{L}_t)$. It's well-known that the spectral gap exists for $\mathcal{L}_t$, if and only if the following Poincare inequality holds

$$Var_{\mu_t}(f) := \int f^2 d\mu_t - \left(\int f d\mu_t\right)^2 \leq C_{P_t}\mathcal{E}_t(f) \tag{3.3}$$

where $C_{P_t}^{-1} := \lambda_t$ is the spectral gap, which is defined as

$$\lambda_t := \inf\left\{\frac{\mathcal{E}_t(g)}{\|g\|^2_{L^2(\mu_t)}} : g \in C^2, g \neq 0, \int g = 0\right\} . \tag{3.4}$$

A defective logarithmic Sobolev inequality (DLSI) is satisfied if

$$Ent_{\mu_t}\left(f^2\right) := \int f^2 \log f^2 d\mu_t - \int f^2 d\mu_t \log\left(\int f^2 d\mu_t\right) \leq \nu_{\mu_t}\mathcal{E}_t(f) + D_{LS_t}\int f^2 d\mu_t.$$

The logarithmic sobolev inequality (LSI) holds when $D_{LS_t} = 0$ and the smallest $\nu_{\mu_t}$ for which this inequality holds is said to be the log-Sobolev constant for $\mu_t$. It's worth mentioning that from the Rothaus's lemma (See Theorem 5.1.4 in Bakry et al. (2014)), if the spectral gap exists for $\mathcal{L}_t$, a defective LSI, can be transformed into a LSI (see Barthe and Kolesnikov (2008) and references therein).

### 3.2 Global Linear Convergence via Uniform Log-Sobolev Inequalities

Now, we are ready to present the following proposition, which obtains a linear convergence of $Q(\rho_t)$ to its global optima $Q(\rho_*)$ under the condition that the logarithmic sobolev inequality is satisfied uniformly for a family of Gibbs measures satisfying $\tilde{\rho}_t \propto \exp(-\lambda^{-1}U(\theta, \rho_t))$.

**Proposition 2.** *The functional $Q(\rho)$ has a unique global optimal solution $\rho_*$. Furthermore, when the logarithmic sobolev inequality is satisfied uniformly with a coefficient $\nu > 0$ for all measures $\tilde{\rho}_t \propto \exp(-\lambda^{-1}U(\theta, \rho_t))$ with $t \geq 0$, then $Q(\rho_t)$ is monotonically non-increasing with $\rho_t$ converges weakly to $\rho_*$. Furthermore, we have for $t \geq 0$ and any $\rho_0$ with $Q(\rho_0) < \infty$,*

$$Q(\rho_t) - Q(\rho_*) \leq [Q(\rho_0) - Q(\rho_*)] * \exp(-2\lambda\nu^{-1}t) \tag{3.5}$$

*where $\nu > 0$ is a problem dependent constant.*

*Proof.* Since the functional $Q(\rho)$ is convex with respect to $\rho$, it follows from standard arguments in the literature that there is a unique global optimal solution $\rho_*$ to $Q(\rho)$ which $\rho_t$ weakly converges to (Mei et al. (2018)).

By the decreasing lemma, we have

$$\frac{dQ(\rho_t)}{dt} = -\int_{\mathbb{R}^{d+1}} \|\nabla_\theta g(\theta, \rho_t)\|_2^2 \rho_t(\theta) d\theta = -\lambda^2 \int_{\mathbb{R}^{d+1}} \|\nabla_\theta \log\exp(\lambda^{-1} g(\theta, \rho_t))\|_2^2 \rho_t(\theta) d\theta . \tag{3.6}$$

If the logarithmic sobolev inequality is satisfied uniformly with a problem dependent coefficient $\nu > 0$ for all measures $\tilde{\rho}_t \propto \exp(-\lambda^{-1}U(\theta, \rho_t)) = \frac{\rho_t(\theta)\exp(-\lambda^{-1}g(\rho_t,\theta))}{\mathbb{E}_{\theta\sim\rho_t}\exp(-\lambda^{-1}g(\rho_t,\theta))}$. Therefore, we obtain for all

$t \geq 0$, $Ent_{\tilde{\rho}_t}\left(f^2\right) \leq \nu \mathcal{E}_t(f)$ . If we set $f(\theta) = \sqrt{\rho_t(\theta)/\tilde{\rho}_t(\theta)} \in L^2(\tilde{\rho}_t)$, we obtain

$$2\nu^{-1} KL(\rho_t \| \tilde{\rho}_t) \leq \mathbb{E}_{\theta \sim \rho_t} \left\| \nabla \log \frac{\rho_t(\theta)}{\tilde{\rho}_t(\theta)} \right\|_2^2 . \tag{3.7}$$

Therefore, we obtain

$$\frac{dQ(\rho_t)}{dt} \leq -2\lambda^2 \nu^{-1} \mathbb{E}_{\theta \sim \rho_t} \log \frac{\mathbb{E}_{\theta \sim \rho_t} \exp(-\lambda^{-1} g(\rho_t, \theta))}{\exp(-\lambda^{-1} g(\rho_t, \theta))} . \tag{3.8}$$

Since $\phi(\cdot, \cdot)$ is convex on its first argument, we obtain that $Q(\rho)$ is a strongly convex functional with respect to $\rho$, and it is straightforward to obtain,

$$\begin{aligned} Q(\rho_t) - Q(\rho_*) &\leq \int_{\mathbb{R}^{d+1}} g_t(\theta, \rho_t)(\rho_t(\theta) - \rho_*(\theta)) d\theta - \lambda KL(\rho_* \| \rho_t) \\ &\leq \lambda \log \mathbb{E}_{\theta \sim \rho_t} \exp(-\lambda^{-1} g_t(\theta, \rho_t)) + \mathbb{E}_{\theta \sim \rho_t} g_t(\theta, \rho_t) . \end{aligned} \tag{3.9}$$

It follows that

$$\frac{d[Q(\rho_t) - Q(\rho_*)]}{dt} \leq -2\lambda \nu^{-1} [Q(\rho_t) - Q(\rho_*)] \tag{3.10}$$

which implies the linear convergence of $Q(\rho_t)$ to $Q(\rho_*)$ . □

**Remark 5.** *We shall note that there are two Arxiv paper (Chizat (2022); Nitanda et al. (2022)) that already obtain the above results prior to the submission of this manuscipt. Our work is done independently to them and the difference is that we assume the uniform LSI holds along the trajectory rather than for any probability measure in $\mathcal{P}_2(\mathbb{R}^d)$ as in (Chizat (2022); Nitanda et al. (2022)). Although this weaker assumption is straightforward from a technical perspective, we stress it because it can lead to a more refined time-dependent estimate of LSI constants in Theorem 2.*
*From a higher perspective, the difficulty to obtain the convergence rate arises from the nonlinearity in the drift term $-\nabla_\theta U(\theta_t, \rho_t)$ of the continuous stochastic process corresponds to the Mckean-Vlason PDE, in contrast to the classical linear Fokker-Planck equation that describes the evolution dynamics of Langevin algorithm where the drift term remains unchanged during the evolution and if we put some isoperimetry condition on the stationary distribution determined by the drift term, such as it satisfies the Logarithmic Sobolev Inequality (LSI), then the linear convergence to the stationary process can be obtained and the convergence rate can be controlled by the LSI constant of the stationary distribution (Ma et al. (2019)). To overcome such difficulty, one can simply adapt the proofs in (Ma et al. (2019)) and substitute the stationary process $p^* \propto \exp(-\lambda^{-1} U(\theta))$ by the varying Gibbs distribution $\tilde{\rho}_t \propto \exp(-\lambda^{-1} U(\theta, \rho_t))$. Then the convergence rate can be controlled by the uniform Log-Sobolev constant of $\tilde{\rho}_t$ along the optimization trajectory.*

**Remark 6.** *In a broader context, the idea of assuming uniform LSI to obtain the linear convergence of McKean-Vlasov process already exists in the literature, for example, in Yang et al. (2020) and Guillin et al. (2019). The fundamental difference between these works and our contribution is that these works directly assume the uniform LSI holds while we prove that the uniform LSI holds in the context of vanilla two-layer neural networks.*

# 4 Estimation of the Uniform Log-Sobolev Constant

In the previous section, we have obtained the linear convergence rate for noisy SGD in training two-layer neural networks in the mean-field regime. But the whole story is far from complete, since this linear convergence result relies on the condition that $\tilde{\rho}_t \propto \exp(-\lambda^{-1}U(\theta, \rho_t))$ satisfies the LSI with coefficients $\nu_t$ and there exist a uniform upper bound $\nu > 0$ on $\nu_t$ for all $t \geq 0$. Hence, in the next, we need to estimate the upper bound of the log-Sobolev constants for $\tilde{\rho}_t \propto \exp(-\lambda^{-1}U(\theta, \rho_t))$.

There are many literatures on establishing the logarithmic sobolev constant for Gibbs distributions of the form $\mu_t \propto e^{-U_t}$. The most simple and fundamental one is the Bakery-Emery Criteria (Bakry and Émery (1985)): if $U_t$ is $C^2$ and strongly convex, i.e., $\nabla^2 U_t \succeq KI$ for some $K > 0$, then $\mu_t$ satisfies the LSI with $\nu_{\mu_t} = 1/K$. However, in many cases, $U_t$ is nonconvex.

In nonconvex cases, the standard tool to deduce the the LSI is the Holley-Stroock perturbation principle (Holley and Stroock (1986)). The idea is that if we can find a strongly convex function $\hat{U}_t$ that is a bounded perturbation of $U_t$, then $U_t$ satisfies the LSI with a worse logarithmic Sobolev constant than $\hat{U}_t$. From an intuitive understanding, in this case, $U_t$ should still be dominated by a strongly convex function, plus some minor bounded nonconvex function as perturbations. This method is adopted in some machine learning literatures of analyzing the Langevin sampling with nonconvex potentials, for example, see Ma et al. (2019) , where $U_t$ is strongly convex outside a ball with fixed radius. In this case, we can construct a globally strongly convex function $\hat{U}_t$ that is the bounded perturbation of $U_t$, hence deducing the LSI.

There are some other extensions to obtain the LSI beyond Bakery-Emery criteria. For example, if there is a constant lower bound on the Hessian, i.e., $\nabla^2 U_t \succeq -KI$ for some $K > 0$, and some Lyapunov function criteria is satisfied, then we can also obtain the LSI for $\mu_t$. This method also appears in analyzing the global convergence of SGLD, for example, see Raginsky et al. (2017).

There are some recent works that establish the logarithmic sobolev inequality when the Bakery-Emery Curvature is not lower bounded, which is considered in Barthe and Kolesnikov (2008) and later Cattiaux et al. (2010) extends the results of Barthe and Kolesnikov (2008) in Euclidean space to the general Riemann Manifold. To our best knowledge, this approach has not been considered in any machine learning literature.

However, it is generally hard to estimate the log-Sobolev constants in the absence of convexity, which is the case when $U_t$ is a nonconvex potential for optimizing a two-layer infinite-width neural network. As will shown in the later analysis, the key technical difficulty is that the lower bound of the Hessian of the two-layer infinite-width neural network grows linearly in terms of the norm of the neuron. How can we overcome such difficulty? Notice that we have adopted a regularizer $r(\theta)$: if the Hessian $\nabla^2 r(\theta)$ grows sufficiently fast as a function of $\|\theta\|_2$, then it may help us control the Hessian of $U_t$. Observe that we cannot adopt the original weight decay $r(\theta) = \frac{\beta}{2}\|\theta\|_2^2$, since the Hessian $\nabla^2 r(\theta) = \beta I$, which is a constant and cannot control the Hessian of the whole $U_t$. This motivates us to consider higher order regularizers beyond weight decay.

In the next two subsections, we consider two cases of $r(\theta)$ where we can bound the log-Sobolev constants uniformly for all $\mu_t \propto e^{-U_t}$.

## 4.1 Estimate of the Logarithmic Sobolev Constant via Holley-Stroock Perturbation Principle

First, we consider a special case $r(\theta) = \beta\|\theta\|_2^4$, which can be regarded as a variant of weight decay by enhancing more decay on weights with larger Euclidean norm. In such case, $U(\theta, \rho_t)$ will

be strongly convex outside a ball with some radius $R_{m,\beta} > 0$ and locally smooth on any compact set in $\mathbb{R}^d$. Using similar construction as Ma et al. (2019) and then adopting the Holley-Stroock perturbation principle, we can bound the LSI constant uniformly for all $\mu_t \propto e^{-U_t}$. As will seen in the proof, our analysis can be directly extended to any $r(\theta)$ in $C^2$ that grows not slower than $\|\theta\|_2^4$.

**Theorem 1** (Proved in Appendix A)**.** *When $r(\theta) = \beta\|\theta\|_2^4$, we have the following uniform upper bound on the Logarithmic Sobolev constant $\nu_{\rho_t}$,*

$$\sup_{t\geq 0} \nu_{\tilde{\rho}_t} \leq \frac{2}{m\lambda^{-1}} e^{16\lambda^{-1}L_{m,\beta}R_{m,\beta}^2} .$$

*where $m > 0$ can be picked arbitrarily and $R_{m,\beta}, L_{m,\beta} > 0$ are all problem dependent constants (detailed in appendix).*

### 4.2 Time-dependent Estimate of the Logarithmic Sobolev Constant

Notice that when we adopt the Holley-Stroock perturbation principle, the obtained LSI could be too loose since there is a term which grows exponentially with the perturbation. This estimate of Log Sobolev does not depend on the trajectory dynamics, which is also a weakness in existing works (Chizat (2022); Nitanda et al. (2022)). Furthermore, the condition that regularizer $r(\theta)$ grows at least $\|\theta\|_2^4$ fast could be too strong. This motivates us to seek alternative and more general approaches to provide a more refined estimate on the LSI constants: we will show, under some Lyapunov conditions together with a uniform spectral gap, the LSI holds uniformly for all $\mu_t \propto e^{-U_t}$ for any $r(\theta)$ in $C^2$ that grows at least as fast as $\|\theta\|_2^r$ for any $r \geq 3$. Notice that in this case the Bakery-Emery Curvature for $\mu_t$ is not lower bounded, hence we adopt the techniques in Barthe and Kolesnikov (2008); Cattiaux et al. (2010).

We need to put some regularity conditions on the regularizer $r(\theta)$, that is critical in establishing the Lyapnov conditions for the Logarithmic Sobolev ineuqality.

**Assumption 3.** *The regularizer $r(\theta)$ should satisfy*

- *For some $m > 0$, $b \geq 0$ and $p \geq 1$,*

$$\langle \theta, \nabla r(\theta)\rangle \geq m\|\theta\|_2^{2+p} - b, \quad \forall \theta \in \mathbb{R}^{d+1} . \tag{4.1}$$

- *For some $D_1, D_2, D_3, D_4, D_5, D_6, D_7, D_8 > 0$, and $k \geq 3$*

$$(D_1\|\theta\|_2 + D_2)I \preceq \nabla^2 r(\theta) \preceq (D_3\|\theta\|_2^p + D_4)I, \tag{4.2}$$

*and*

$$\|\nabla r(\theta)\|_2 \geq D_5\|\theta\|_2^2 + D_6, \tag{4.3}$$

*and*

$$|r(\theta)| \leq D_7\|\theta\|_2^k + D_8 . \tag{4.4}$$

**Remark 7.** *The above assumptions naturally apply to many natural regularizers such as $r(\theta) = \frac{\beta_1}{2}\|\theta\|_2^2 + \frac{\beta_2}{3}\|\theta\|_2^3$ , which also appears in the analysis of Allen-Zhu and Li (2020).*

As will be shown in the proof, we will first establish a time-dependent estimate of the Logarithmic Sobolev constants $\nu_{\rho_t}$ for a family of Gibbs distributions with potential $U(\theta;\rho_t)$. Then, for $t \geq 0$, we have the following uniform upper bound for $\nu_{\rho_t}$.

**Theorem 2** (Proved in Appendix B)**.** *We have the following uniform upper bound on the Logarithmic Sobolev constant $\nu_{\rho_t}$, for any $r(\theta)$ satisfying Assumption 3:*

$$\sup_{t\geq 0} \nu_{\tilde{\rho}_t} \leq 2\sqrt{\frac{2}{c_1\lambda^{-1}\Phi(0)}} + \frac{1}{c_1} + \left[2\sqrt{\frac{2}{c_1\lambda^{-1}\Phi(0)}}c_2 + \frac{2c_2}{c_1} + 2\right] \inf_{r>\sqrt{\frac{c_2}{c_1\lambda^{-1}\Phi(0)}}} \frac{1 + Cc_1 r^4\lambda^{-1}\Phi(0)e^{2poly(r)+\gamma r^2/2}}{c_1 r^2\lambda^{-1}\Phi(0) - c_2}$$

*where $\gamma, \Phi(0), c_1, c_2 > 0$ are all problem dependent constants (detailed in appendix) and*

$$poly(r) = C_7 r^k + L_1 C_1 r^2 + L_1 C_2 r + C_8 \ .$$

## 5 Conclusion

We have shown that under properly chosen regularizer, the linear convergence to the global optima for noisy SGD can be obtained in training two-layer neural networks in the mean-field regime. The key step is to establish the logarithmic sobolev inequality uniformly for all Gibbs measures determined by the evolving distribution $\rho_t$. From a technical perspective, our theory implies that some growth condition on the regularizer $r(\theta)$ is sufficient for the global linear convergence, where at least cubic growth of $r(\theta)$ is required. Our theory also shows that the log-Sobolev constants for $\mu_t \propto e^{-U_t}$ play a crucial role in determining the convergence rate and different regularizers $r(\theta)$ produce different estimates of the log-Sobolev constants. Our current theory also has limitations and leads to some questions that worth future exploration:

- In practice, we usually let $r(\theta) = \beta\|\theta\|_2^2$ that grows quadratically as done in the original weight decay. Is the cubic growth condition for $r(\theta)$ a necessary condition for the global linear convergence result? To our best knowledge, there is no convergence rate proofs for the case $r(\theta) = \beta\|\theta\|_2^2$ even for a simple vanilla two-layer neural network.
- The main focus of this paper is on analyzing the convergence property of the continuous-time noisy SGD. In practice, we need to adopt discretization, which produces discretization errors. It leaves as our future work on how to control the discretization error carefully to keep the linear convergence property as in continuous time.

## References


Allen-Zhu, Z. and Li, Y. (2020). Feature purification: How adversarial training performs robust deep learning. *arXiv preprint arXiv:2005.10190*.

Allen-Zhu, Z., Li, Y., and Liang, Y. (2018). Learning and generalization in overparameterized neural networks, going beyond two layers. *arXiv preprint arXiv:1811.04918*.

Allen-Zhu, Z., Li, Y., and Song, Z. (2019). A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pages 242–252.

Ambrosio, L., Gigli, N., and Savaré, G. (2008). *Gradient flows: in metric spaces and in the space of probability measures*. Springer Science & Business Media.

Ambrosio, L. and Savaré, G. (2007). Gradient flows of probability measures. *Handbook of differential equations: evolutionary equations*, 3:1–136.

Arora, S., Du, S. S., Hu, W., Li, Z., and Wang, R. (2019). Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. *arXiv preprint arXiv:1901.08584*.

Bakry, D., Barthe, F., Cattiaux, P., and Guillin, A. (2008). A simple proof of the poincaré inequality for a large class of probability measures. *Electronic Communications in Probability*, 13:60–66.

Bakry, D. and Émery, M. (1985). Diffusions hypercontractives. In *Seminaire de probabilités XIX 1983/84*, pages 177–206. Springer.

Bakry, D., Gentil, I., Ledoux, M., et al. (2014). *Analysis and geometry of Markov diffusion operators*, volume 103. Springer.

Barthe, F. and Kolesnikov, A. V. (2008). Mass transport and variants of the logarithmic sobolev inequality. *Journal of Geometric Analysis*, 18(4):921–979.

Cattiaux, P., Guillin, A., Wang, F.-Y., and Wu, L. (2009). Lyapunov conditions for super poincaré inequalities. *Journal of Functional Analysis*, 256(6):1821–1841.

Cattiaux, P., Guillin, A., and Wu, L.-M. (2010). A note on talagrand's transportation inequality and logarithmic sobolev inequality. *Probability theory and related fields*, 148(1-2):285–304.

Chen, Z., Vanden-Eijnden, E., and Bruna, J. (2022). On feature learning in neural networks with global convergence guarantees. *arXiv preprint arXiv:2204.10782*.

Chewi, S., Gouic, T. L., Lu, C., Maunu, T., and Rigollet, P. (2020). Svgd as a kernelized wasserstein gradient flow of the chi-squared divergence. *arXiv preprint arXiv:2006.02509*.

Chizat, L. (2022). Mean-field langevin dynamics: Exponential convergence and annealing. *arXiv preprint arXiv:2202.01009*.

Chizat, L. and Bach, F. (2018). On the global convergence of gradient descent for over-parameterized models using optimal transport. *Advances in neural information processing systems*, 31.

Cordero-Erausquin, D., McCann, R. J., and Schmuckenschläger, M. (2001). A riemannian interpolation inequality à la borell, brascamp and lieb. *Inventiones mathematicae*, 146(2):219–257.

Cordero-Erausquin, D., McCann, R. J., and Schmuckenschläger, M. (2006). Prékopa–leindler type inequalities on riemannian manifolds, jacobi fields, and optimal transport. In *Annales de la Faculté des sciences de Toulouse: Mathématiques*, volume 15, pages 613–635.

Du, S., Lee, J., Li, H., Wang, L., and Zhai, X. (2019). Gradient descent finds global minima of deep neural networks. In *International Conference on Machine Learning*, pages 1675–1685.

Du, S. S., Zhai, X., Poczos, B., and Singh, A. (2018). Gradient descent provably optimizes over-parameterized neural networks.

Fang, C., Lee, J., Yang, P., and Zhang, T. (2021). Modeling from features: a mean-field framework for over-parameterized deep neural networks. In *Conference on learning theory*, pages 1887–1936. PMLR.

Guillin, A., Liu, W., Wu, L., and Zhang, C. (2019). Uniform poincar {\'e} and logarithmic sobolev inequalities for mean field particles systems. *arXiv preprint arXiv:1909.07051*.

Holley, R. and Stroock, D. W. (1986). Logarithmic sobolev inequalities and stochastic ising models.

Hu, K., Ren, Z., Siska, D., and Szpruch, L. (2019). Mean-field langevin dynamics and energy landscape of neural networks. *arXiv preprint arXiv:1905.07769*.

Jacot, A., Gabriel, F., and Hongler, C. (2018). Neural tangent kernel: Convergence and generalization in neural networks. In *Advances in neural information processing systems*, pages 8571–8580.

Johnson, R. and Zhang, T. (2019). A framework of composite functional gradient methods for generative adversarial models. *IEEE transactions on pattern analysis and machine intelligence*.

Lee, J., Bahri, Y., Novak, R., Schoenholz, S. S., Pennington, J., and Sohl-Dickstein, J. (2018). Deep neural networks as gaussian processes.

Li, Y. and Liang, Y. (2018). Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems*, pages 8157–8166.

Li, Y., Ma, T., and Zhang, H. R. (2020). Learning over-parametrized two-layer neural networks beyond ntk. In *Conference on Learning Theory*, pages 2613–2682. PMLR.

Liu, Q. (2017). Stein variational gradient descent as gradient flow. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 3118–3126.

Liu, Q. and Wang, D. (2016). Stein variational gradient descent: a general purpose bayesian inference algorithm. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 2378–2386.

Ma, Y.-A., Chen, Y., Jin, C., Flammarion, N., and Jordan, M. I. (2019). Sampling can be faster than optimization. *Proceedings of the National Academy of Sciences*, 116(42):20881–20885.

Mei, S., Misiakiewicz, T., and Montanari, A. (2019). Mean-field theory of two-layers neural networks: dimension-free bounds and kernel limit. In *Conference on Learning Theory*, pages 2388–2464. PMLR.

Mei, S., Montanari, A., and Nguyen, P.-M. (2018). A mean field view of the landscape of two-layer neural networks. *Proceedings of the National Academy of Sciences*, 115(33):E7665–E7671.

Neal, R. M. (1996). Priors for infinite networks. In *Bayesian Learning for Neural Networks*, pages 29–53. Springer.

Neal, R. M. (2012). *Bayesian learning for neural networks*, volume 118. Springer Science & Business Media.

Nguyen, P.-M. and Pham, H. T. (2020). A rigorous framework for the mean field limit of multilayer neural networks. *arXiv preprint arXiv:2001.11443*.

Nitanda, A. and Suzuki, T. (2020). Functional gradient boosting for learning residual-like networks with statistical guarantees. In *International Conference on Artificial Intelligence and Statistics*, pages 2981–2991. PMLR.

Nitanda, A., Wu, D., and Suzuki, T. (2022). Convex analysis of the mean field langevin dynamics. *arXiv preprint arXiv:2201.10469*.

Raginsky, M., Rakhlin, A., and Telgarsky, M. (2017). Non-convex learning via stochastic gradient langevin dynamics: a nonasymptotic analysis. In *Conference on Learning Theory*, pages 1674–1703. PMLR.

Santambrogio, F. (2015). Optimal transport for applied mathematicians. *Birkäuser, NY*, 55(58-63):94.

Sirignano, J. and Spiliopoulos, K. (2019). Mean field analysis of neural networks: A central limit theorem. *Stochastic Processes and their Applications*.

Villani, C. (2009). *Optimal transport: old and new*, volume 338. Springer.

Williams, C. K. (1997). Computing with infinite networks. In *Advances in neural information processing systems*, pages 295–301.

Winther, O. (2001). Computing with finite and infinite networks. In *Advances in Neural Information Processing Systems*, pages 336–342.

Yang, Z., Zhang, Y., Chen, Y., and Wang, Z. (2020). Variational transport: A convergent particle-basedalgorithm for distributional optimization. *arXiv preprint arXiv:2012.11554*.

Zou, D., Cao, Y., Zhou, D., and Gu, Q. (2018). Stochastic gradient descent optimizes over-parameterized deep relu networks. *arXiv preprint arXiv:1811.08888*.

# A Proofs of Theorem 1

*Proof.* If we expand $U(\theta, \rho_t)$, we get

$$U(\theta, \rho_t) = U(w, u; \rho_t) = \mathbb{E}_{x,y}\phi_1'(f(\rho, x), y)uh(w, x) + \beta\|\theta\|_2^4 \ . \tag{A.1}$$

Then the gradient of $U(\theta, \rho_t)$ becomes

$$\nabla U(w, u; \rho_t) = \mathbb{E}_{x,y}\left\{\phi_1'(f(\rho, x), y)\begin{bmatrix} u\nabla_1 h(w, x) \\ \\ h(w, x) \end{bmatrix}\right\} + 4\beta\|\theta\|_2^2\theta \ , \tag{A.2}$$

and the Hessian of $U(\theta, \rho_t)$ is

$$\nabla^2 U(\theta; \rho_t) = \mathbb{E}_{x,y}\left\{\phi_1'(f(\rho, x), y)\begin{bmatrix} u\nabla_1^2 h(w, x) & \nabla_1 h(w, x) \\ \\ \nabla_1^T h(w, x) & 0 \end{bmatrix}\right\} + 4\beta\|\theta\|_2^2 I + 8\beta\theta\theta^T \ . \tag{A.3}$$

□

Furthermore, we have

$$\begin{aligned}
&|\lambda_{min}(\nabla^2\mathbb{E}_{x,y}\phi_1'(f(\rho, x), y)uh(w, x))|^2 \\
&\leq \left\| \begin{bmatrix} \mathbb{E}_{x,y}\{\phi_1'(f(\rho, x), y)u\nabla_1^2 h(w, x)\} & \mathbb{E}_{x,y}\{\phi_1'(f(\rho, x), y)\nabla_1 h(w, x)\} \\ \\ \mathbb{E}_{x,y}\{\phi_1'(f(\rho, x), y)\nabla_1^T h(w, x)\} & 0 \end{bmatrix} \right\|_F^2 \\
&= \|\mathbb{E}_{x,y}\{\phi_1'(f(\rho, x), y)u\nabla_1^2 h(w, x)\}\|_F^2 + 2\|\mathbb{E}_{x,y}\{\phi_1'(f(\rho, x), y)\nabla_1 h(w, x)\}\|_2^2 \\
&\leq L_1^2\|u\nabla_1^2 h(w, x)\}\|_F^2 + 2L_1^2 C_3^2 \\
&\leq dL_1^2\|\nabla_1^2 h(w, x)\}\|_{op}^2\|\theta\|_2^2 + 2L_1^2 C_3^2 \\
&\leq dL_1^2 C_4^2\|\theta\|_2^2 + 2L_1^2 C_3^2 \ .
\end{aligned} \tag{A.4}$$

Hence, we have

$$\nabla^2 U(\theta; \rho_t) \succeq \left(4\beta\|\theta\|_2^2 - \sqrt{dL_1^2 C_4^2\|\theta\|_2^2 + 2L_1^2 C_3^2}\right) I \succeq \left(4\beta\|\theta\|_2^2 - \sqrt{d}L_1 C_4\|\theta\|_2 - \sqrt{2}L_1 C_3\right) I \tag{A.5}$$

Obviously, we can let

$$R_{m,\beta} = \frac{\sqrt{d}L_1 C_4 + \sqrt{dL_1^2 C_4^2 + 16(m + \sqrt{2}\beta L_1 C_3)}}{8\beta} \ .$$

Then we obtain that $U(\theta; \rho_t)$ is $m$-strongly convex for $\|\theta\|_2 \geq R_{m,\beta}$ . For convenience and clarity of proof, without loss of generality, we impose $h(0, x) = 0$ for any $x$. Then we have $\nabla U(0, \rho_t) = 0$. Furthermore, we have the following bound showing that $U(\theta; \rho_t)$ is local Lipschitz smooth.

$$\|\nabla^2 U(\theta; \rho_t)\|_{op} \leq 12\beta\|\theta\|_2^2 + \sqrt{dL_1^2 C_4^2\|\theta\|_2^2 + 2L_1^2 C_3^2} \ . \tag{A.6}$$

Then, we can let

$$L_{m,\beta} = 48\beta R_{m,\beta}^2 + \sqrt{4dL_1^2C_4^2R_{m,\beta}^2 + 2L_1^2C_3^2} \tag{A.7}$$

and conclude that $U(\theta;\rho_t)$ is $L_{m,\beta}$-Lipschitz smooth for $\|\theta\|_2 \leq 2R_{m,\beta}$ . Then the results can be obtained from almost the same construction and proofs of Lemma 3 in Ma et al. (2019), by modifying the global $L$-Lipschitz smooth assumption to local Lipschitz smooth assumption within the ball $B\left(0, 2R_{m,\beta}\right)$ .

# B Proofs of Theorem 2

The following lemma is essential in establishing the bound of log-Sobolev constants for distributions $\tilde{\rho}_t(\theta)$. The proof of the lemma is nothing unusual but adapting and combining existing results in the literature. More specifically, it can be obtained by combining results from Theorem 1.4 in Bakry et al. (2008), Proposition 4.3 in Cattiaux et al. (2010), and the Rothaus' lemma.

**Lemma 2.** *For a probability distribution $\rho_t = \exp\left(-U_t(\theta)\right)/Z$, where $U \in C^2$ and $Z = \int_{\mathbb{R}^{d+1}} e^{-U_t(\theta)}d\theta < +\infty$. If $U_t(\theta)$ satisfies the following two conditions*

- $\nabla^2 U_t(\theta) \succeq -\Phi(\|\theta\|_2)I$ *where $\Phi$ is some positive non-decreasing continuous function on $\mathbb{R}^+$.*
- *There is some nonnegative and locally Lipschitz function $V : \mathbb{R}^{d+1} \to \mathbb{R}^+ \cup \{0\}$ such that for some constants $c_1, c_2 > 0$,*

$$\mathcal{L}_t V(\theta) + \|\nabla V(\theta)\|^2 \leq -c_1\|\theta\|_2^2\Phi(2\|\theta\|_2) + c_2 \ . \tag{B.1}$$

*Then $\rho_t$ satisfies the log-Sobolev inequality with*

$$\nu_{\rho_t} = 2\sqrt{\frac{2}{c_1\Phi(0)}} + \frac{1}{c_1} + \left[2\sqrt{\frac{2}{c_1\Phi(0)}}c_2 + \frac{2c_2}{c_1} + 2\right] \inf_{r > \sqrt{\frac{c_2}{c_1\Phi(0)}}} \frac{1 + Cc_1\Phi(0)r^4 \sup_{\|\theta\|_2 \leq r} \exp(V(\theta))e^{Osc_r(U_t)}}{c_1r^2\Phi(0) - c_2}$$

*where $C > 0$ is a universal constant and*

$$Osc_r(U_t) = \sup_{\|\theta\|_2 \leq r} U_t(\theta) - \inf_{\|\theta\|_2 \leq r} U_t(\theta) \ .$$

*Proof.* The key step of establishing the LSI in the unbounded curvature case is by exploiting the above-tangent lemma as did in originally in Barthe and Kolesnikov (2008) and later in Cattiaux et al. (2010). The above-tangent lemma consider a general results on a Riemannian manifold (Cordero-Erausquin et al. (2001, 2006)), and we adapted it to the Euclidean space $\mathbb{R}^{d+1}$ (i.e., with zero Ricci curvature):

let $d\mu_t(\theta) = Z^{-1}e^{-U_t(\theta)}d\theta$ be a probability measure on $\mathbb{R}^{d+1}$ and denote by $f \cdot \mu_t$ the measure with density $f$ with respect to $\mu_t$. Consider two compactly supported nonnegative functions $g$ and $h$ such that $\nu := g \cdot \mu_t$ and $h \cdot \mu_t$ are probability measures. Let $T(\theta) = \theta + \nabla\tau(\theta)$ be the optimal transport minimizing the quadratic transport cost and pushing forward $g \cdot \mu_t$ to $h \cdot \mu_t$. Then it

holds

$$Ent_{\mu_t}(g) \leq Ent_{\mu_t}(h) - \int_{\mathbb{R}^{d+1}} \langle \nabla\tau, \nabla g\rangle d\mu_t + \int_{\mathbb{R}^{d+1}} \mathcal{D}_{U_t}(\theta, T(\theta)) g d\mu_t \tag{B.2}$$

where the convexity defect is defined as

$$\mathcal{D}_{U_t}(\theta, T(\theta)) = -\int_0^1 (1-s)\nabla\tau(\theta)^T \nabla^2 U_t(\gamma(s))\nabla\tau(\theta) ds \ , \quad \gamma(s) = \theta + s\nabla\tau(\theta) \ , \quad s \in [0,1].$$

Now, we choose a sequence of $\mu_t$-probability measures $\mu_n = h_n\mu_t$ with $h_n \in C_0^\infty(\mathbb{R}^{d+1})$, such that $W_2(\mu_n, \mu_t) \to 0, H(\mu_n|\mu_t) := Ent_{\mu_t}(h_n) \to 0$ and $I(\mu_n|\mu_t) := \mathcal{E}_t(\sqrt{h_n}) \to 0$ . In the following, we apply the above-tangent lemma to $\nu, \mu_n$.

Hence,

$$\begin{aligned}
&\left| \int_{\mathbb{R}^{d+1}} \langle \nabla\tau, \nabla g\rangle d\mu_t \right| \\
&= \left| \int_{\mathbb{R}^{d+1}} 2\sqrt{g}\langle \nabla\tau, \nabla\sqrt{g}\rangle d\mu_t \right| \\
&\leq 2\sqrt{\int_{\mathbb{R}^{d+1}} \|\nabla\tau\|_2^2 g d\mu_t \int_{\mathbb{R}^{d+1}} \|\nabla\sqrt{g}\|_2^2 d\mu_t} \\
&= 2W_2(\nu, \mu_n)\sqrt{\mathcal{E}_t(\sqrt{g})} \ .
\end{aligned} \tag{B.3}$$

Then we bound the defect of convexity by our condition

$$\begin{aligned}
&\mathcal{D}_{U_t}(\theta, T(\theta)) \\
&= -\int_0^1 (1-s)\nabla\tau(\theta)^T \nabla^2 U_t(\gamma(s))\nabla\tau(\theta) ds \\
&\leq \int_0^1 (1-s)\|\nabla\tau(\theta)\|_2^2 \Phi\left(\|\gamma(s)\|_2\right) ds \ .
\end{aligned} \tag{B.4}$$

By the triangle inequality, we have

$$\|\nabla\tau(\theta)\|_2 = \|T(\theta) - \theta\|_2 \leq 2\max\{\|T(\theta)\|_2, \|\theta\|_2\} \tag{B.5}$$

and

$$\|\gamma(s)\|_2 = \|\theta + s\nabla\tau(\theta)\|_2 = \|\theta + s(T(\theta) - \theta)\|_2 \leq 2\max\{\|\theta\|_2, \|T(\theta)\|_2\} \tag{B.6}$$

Therefore we can bound the last term as follows:

$$\begin{aligned}
&\int_{\mathbb{R}^{d+1}} \mathcal{D}_{U_t}(\theta, T(\theta)) g d\mu_t \\
&\leq 2\int_{\mathbb{R}^{d+1}} \max\{\|T(\theta)\|_2^2, \|\theta\|_2^2\} \Phi\left(2\max\{\|\theta\|_2, \|T(\theta)\|_2\}\right) g d\mu_t \\
&\leq 2\left( \int_{\mathbb{R}^{d+1}} \|T(\theta)\|_2^2 \Phi\left(2\|T(\theta)\|_2\right) g d\mu_t + \int_{\mathbb{R}^{d+1}} \|\theta\|_2^2 \Phi\left(2\|\theta\|_2\right) g d\mu_t \right)
\end{aligned} \tag{B.7}$$

By Lemma 3, we have

$$\int_{\mathbb{R}^{d+1}} \|T(\theta)\|_2^2 \Phi\left(2\|T(\theta)\|_2\right) g d\mu_t \leq \frac{1}{c_1}(I(\mu_n|\mu_t) + c_2) \tag{B.8}$$

and

$$\int_{\mathbb{R}^{d+1}} \|\theta\|_2^2 \Phi\left(2\|\theta\|_2\right) g d\mu_t \leq \frac{1}{c_1}(\mathcal{E}_t(\sqrt{g}) + c_2) \ . \tag{B.9}$$

Combining the above results and letting $n \to \infty$, we obtain

$$Ent_{\mu_t}(g) \leq 2W_2(\nu,\mu)\sqrt{\mathcal{E}_t(\sqrt{g})} + \frac{1}{c_1}(\mathcal{E}_t(\sqrt{g}) + 2c_2) \ . \tag{B.10}$$

Similarly, we also have

$$W_2(\nu,\mu) \leq 2\left(\int_{\mathbb{R}^{d+1}} \|\theta\|_2^2 g d\mu_t + \int_{\mathbb{R}^{d+1}} \|\theta\|_2^2 d\mu_t\right) \leq \frac{2}{c_1\Phi(0)}(\mathcal{E}_t(\sqrt{g}) + 2c_2) \ . \tag{B.11}$$

Therefore, we obtain the defective LSI,

$$Ent_{\mu_t}(g) \leq \left(2\sqrt{\frac{2}{c_1\Phi(0)}} + \frac{1}{B}\right)\mathcal{E}_t(\sqrt{g}) + 2c_2\left(\sqrt{\frac{2}{c_1\Phi(0)}} + \frac{1}{B}\right) \ . \tag{B.12}$$

Then, we need to estimate the spectral gap under the Lyapunov condition (B.1).
If we let $V(\theta) = \log W(\theta)$ for some $W : \mathbb{R}^{d+1} \to [1,\infty)$, then with some simple algebra we obtain

$$\frac{\mathcal{L}_t W(\theta)}{W(\theta)} \leq -c_1\|\theta\|_2^2\Phi(2\|\theta\|_2) + c_2 \leq -c_1\|\theta\|_2^2\Phi(0) + c_2 \ . \tag{B.13}$$

When $\theta \in B^c(0,r)$ and choose $r \geq \sqrt{\frac{c_2}{c_2\Phi(0)}}$, we have

$$\frac{\mathcal{L}_t W(\theta)}{W(\theta)} \leq -(c_1 r^2\Phi(0) - c_2) \ . \tag{B.14}$$

When $\theta \in B(0,r)$, we have

$$\frac{\mathcal{L}_t W(\theta)}{W(\theta)} \leq c_2 \ . \tag{B.15}$$

Hence,

$$\frac{\mathcal{L}_t W(\theta)}{W(\theta)} \leq -(c_1 r^2\Phi(0) - c_2) + c_1 r^2\Phi(0)\mathbb{1}_{B(0,r)}(\theta) \ . \tag{B.16}$$

Then,

$$\mathcal{L}_t W(\theta) \leq -(c_1 r^2\Phi(0) - c_2)W(\theta) + c_1 r^2\Phi(0) \sup_{\|\theta\|_2 \leq r} \exp\left(V(\theta)\right)\mathbb{1}_{B(0,r)}(\theta) \ . \tag{B.17}$$

By Theorem 1.4 in Bakry et al. (2008), and a rough estimate of the Poincare constant for the truncated Gibbs measure by $Cr^2 e^{Osc_r(U_t)}$, we obtain that $\mu_t$ also satisfies the Poincare inequality with constant

$$C_{P_t} = \inf_{r > \sqrt{\frac{c_2}{c_1 \Phi(0)}}} \frac{1 + C c_1 \Phi(0) r^4 \sup_{\|\theta\|_2 \le r} \exp(V(\theta)) e^{Osc_r(U_t)}}{c_1 r^2 \Phi(0) - c_2} \ . \tag{B.18}$$

By the Rothaus's lemma, the DLSI together with the spectral gap of $\mathcal{L}_t$ implies the LSI with constant

$$\begin{aligned}
\nu_{\rho_t} &= 2\sqrt{\frac{2}{c_1\Phi(0)}} + \frac{1}{c_1} + \left[2\sqrt{\frac{2}{c_1\Phi(0)}}c_2 + \frac{2c_2}{c_1} + 2\right] C_{P_t} \\
&= 2\sqrt{\frac{2}{c_1\Phi(0)}} + \frac{1}{c_1} + \left[2\sqrt{\frac{2}{c_1\Phi(0)}}c_2 + \frac{2c_2}{c_1} + 2\right] \inf_{r > \sqrt{\frac{c_2}{c_1 \Phi(0)}}} \frac{1 + C c_1 \Phi(0) r^4 \sup_{\|\theta\|_2 \le r} \exp(V(\theta)) e^{Osc_r(U_t)}}{c_1 r^2 \Phi(0) - c_2} \ .
\end{aligned}$$

□

Now, we are ready to establish a time-dependent estimate of the Logarithmic Sobolev Constants for a family of Gibbs distributions with potential $U(\theta; \rho_t)$.

**Theorem 3.** *For $\tilde{\rho}_t(\theta) \propto \exp\left(-\lambda^{-1} U(\theta, \rho_t)\right)$, the logarithmic Sobolev Contant $\nu_{\rho_t}$ satisfies*

$$\nu_{\tilde{\rho}_t} = 2\sqrt{\frac{2}{c_1\lambda^{-1}\Phi(0)}} + \frac{1}{c_1} + \left[2\sqrt{\frac{2}{c_1\lambda^{-1}\Phi(0)}}c_2 + \frac{2c_2}{c_1} + 2\right] \inf_{r > \sqrt{\frac{c_2}{c_1\lambda^{-1}\Phi(0)}}} \frac{1 + C c_1 r^4 \lambda^{-1} \Phi(0) e^{Osc_r(U_t) + \gamma r^2/2}}{c_1 r^2 \lambda^{-1} \Phi(0) - c_2}$$

*where $\Phi(0) = \sqrt{4L_1^2 C_3^2 + 2(d+1)D_4^2}$; $c_1 = 1$; and*

$$\begin{aligned}
c_2 = c_2(R) &:= c_1 R^2 \Phi(2R) + \gamma(d+1) + \gamma(\gamma + L_1 C_1 + L_1 C_3^2) R^2 + \gamma L_1 C_2 R + \gamma b \\
&= (2R)^{p/2} \sqrt{2(d+1)D_3^2} + (2R)\sqrt{2dL_1^2 C_4^2} + 2(2R)^{p/4}\sqrt{(d+1)D_3 D_4} \\
&+ R^2 \sqrt{4L_1^2 C_3^2 + 2(d+1)D_4^2} + \gamma(d+1) + \gamma(\gamma + L_1 C_1 + L_1 C_3^2) R^2 + \gamma L_1 C_2 R + \gamma b
\end{aligned}$$

*and when $p = 1$, we can set any*

$$\begin{aligned}
R^2 \ge \max\Bigg\{ & 1, \left(\frac{\sqrt{4L_1^2 C_3^2 + 2(d+1)D_4^2} + \gamma(\gamma + L_1 C_1 + L_1 C_3^2)}{m}\right)^{\frac{2}{p}}, \\
& \left(\frac{\gamma L_1 C_2}{m}\right)^{\frac{2}{1+p}}, \left(\frac{\gamma(b+d+1)}{m}\right)^{\frac{2}{2+p}}, \left(\frac{2\sqrt{(d+1)D_3 D_4}}{m}\right)^{\frac{4}{p}} \Bigg\}
\end{aligned}$$

*and let*

$$\gamma = \frac{1}{m} 2^{p/2} \left(\sqrt{2(d+1)} D_3 + 2L_1 C_4 \sqrt{d}\right) + 4 \ ; \tag{B.19}$$

*when $p > 1$, we can set any*

$$R^2 \geq \max\left\{1, \left(\frac{\sqrt{4L_1^2C_3^2 + 2(d+1)D_4^2} + \gamma(\gamma + L_1C_1 + L_1C_3^2)}{m}\right)^{\frac{2}{p}}, \right.$$
$$\left.\left(\frac{\gamma L_1 C_2}{m}\right)^{\frac{2}{1+p}}, \left(\frac{\gamma(b+d+1)}{m}\right)^{\frac{2}{2+p}}, \left(\frac{2L_1C_4\sqrt{d}}{m}\right)^{\frac{2}{p-1}}, \left(\frac{2\sqrt{(d+1)D_3D_4}}{m}\right)^{\frac{4}{p}}\right\}$$

*and let*

$$\gamma = \frac{1}{m}2^{p/2}\sqrt{2(d+1)}D_3 + 5 \ . \tag{B.20}$$

*Proof.* If we expand $U(\theta, \rho_t)$, we get

$$U(\theta, \rho_t) = U(w, u; \rho_t) = \mathbb{E}_{x,y}\phi_1'(f(\rho, x), y)uh(w, x) + r(\theta) \ . \tag{B.21}$$

Then the gradient of $U(\theta, \rho_t)$ becomes

$$\nabla U(w, u; \rho_t) = \mathbb{E}_{x,y}\left\{\phi_1'(f(\rho, x), y)\begin{bmatrix} u\nabla_1 h(w, x) \\ \\ h(w, x) \end{bmatrix}\right\} + \nabla r(\theta) \ , \tag{B.22}$$

and the Hessian of $U(\theta, \rho_t)$ is

$$\nabla^2 U(\theta; \rho_t) = \mathbb{E}_{x,y}\left\{\phi_1'(f(\rho, x), y)\begin{bmatrix} u\nabla_1^2 h(w, x) & \nabla_1 h(w, x) \\ \\ \nabla_1^T h(w, x) & 0 \end{bmatrix}\right\} + \nabla^2 r(\theta) \ . \tag{B.23}$$

Then, we have the following bound on $\nabla U(w, u; \rho_t)$,

$$\begin{aligned}
&\|\nabla U(w, u; \rho_t)\|_2 \\
&= \left\|\mathbb{E}_{x,y}\left\{\phi_1'(f(\rho, x), y)\begin{bmatrix} u\nabla_1 h(w, x) \\ \\ h(w, x) \end{bmatrix}\right\} + \nabla r(\theta)\right\|_2 \\
&\geq \|\nabla r(\theta)\|_2 - \left\|\mathbb{E}_{x,y}\left\{\phi_1'(f(\rho, x), y)\begin{bmatrix} u\nabla_1 h(w, x) \\ \\ h(w, x) \end{bmatrix}\right\}\right\|_2 \\
&\geq \left(D_5\|\theta\|_2^2 + D_6\right) - L_1\left(|u|\|\nabla_1 h(w, x)\|_2 + |h(w, x)|\right) \\
&\geq \left(D_5\|\theta\|_2^2 + D_6\right) - L_1\left(|u|C_3 + C_1\|w\|_2 + C_2\right) \\
&\geq \left(D_5\|\theta\|_2^2 + D_6\right) - L_1\left(\|\theta\|_2 C_3 + C_1\|\theta\|_2 + C_2\right) \\
&= D_5\|\theta\|_2^2 - (L_1C_3 + C_1)\|\theta\|_2 + (D_6 - L_1C_2) \ .
\end{aligned} \tag{B.24}$$

Besides, we can bound the smallest eigenvalue of $\nabla^2 U(\theta; \rho_t)$ as follows,

$$|\lambda_{min}(\nabla^2 U(\theta; \rho_t)|^2$$

$$
\begin{aligned}
&\leq trace(\nabla^2 U(\theta;\rho_t)\nabla^2 U(\theta;\rho_t)^T) \\
&= \underbrace{\left\| \begin{bmatrix} \mathbb{E}_{x,y}\{\phi_1'(f(\rho,x),y)u\nabla_1^2 h(w,x)\} & \mathbb{E}_{x,y}\{\phi_1'(f(\rho,x),y)\nabla_1 h(w,x)\} \\ \mathbb{E}_{x,y}\{\phi_1'(f(\rho,x),y)\nabla_1^T h(w,x)\} & 0 \end{bmatrix} \right\|_F^2}_{T_1} \\
&+ \underbrace{2\, trace\left( \begin{bmatrix} \mathbb{E}_{x,y}\{\phi_1'(f(\rho,x),y)u\nabla_1^2 h(w,x)\} & \mathbb{E}_{x,y}\{\phi_1'(f(\rho,x),y)\nabla_1 h(w,x)\} \\ \mathbb{E}_{x,y}\{\phi_1'(f(\rho,x),y)\nabla_1^T h(w,x)\} & 0 \end{bmatrix} \nabla^2 r(\theta)^T \right)}_{T_2} \\
&+ \underbrace{trace(\nabla^2 r(\theta)\nabla^2 r(\theta)^T)}_{T_3} \ . \qquad \text{(B.25)}
\end{aligned}
$$

To bound $T_1$, we have

$$
\begin{aligned}
T_1 &= \|\mathbb{E}_{x,y}\{\phi_1'(f(\rho,x),y)u\nabla_1^2 h(w,x)\}\|_F^2 + 2\|\mathbb{E}_{x,y}\{\phi_1'(f(\rho,x),y)\nabla_1 h(w,x)\}\|_2^2 \\
&\leq L_1^2\|u\nabla_1^2 h(w,x)\}\|_F^2 + 2L_1^2 C_3^2 \\
&\leq dL_1^2\|\nabla_1^2 h(w,x)\}\|_{op}^2\|\theta\|_2^2 + 2L_1^2 C_3^2 \\
&\leq dL_1^2 C_4^2\|\theta\|_2^2 + 2L_1^2 C_3^2 \ . \qquad \text{(B.26)}
\end{aligned}
$$

Similarly, we can bound $T_3$ as follows

$$
\begin{aligned}
T_3 &= trace(\nabla^2 r(\theta)\nabla^2 r(\theta)^T) \\
&\leq (D_3\|\theta\|_2^p + D_4)^2 trace(I) \\
&= (d+1)D_3^2\|\theta\|_2^{2p} + 2(d+1)D_3D_4\|\theta\|_2^p + (d+1)D_4^2 \ . \qquad \text{(B.27)}
\end{aligned}
$$

Finally, we get the bound of $T_2$ as follows

$$
\begin{aligned}
T_2 &\leq \left\| \begin{bmatrix} \mathbb{E}_{x,y}\{\phi_1'(f(\rho,x),y)u\nabla_1^2 h(w,x)\} & \mathbb{E}_{x,y}\{\phi_1'(f(\rho,x),y)\nabla_1 h(w,x)\} \\ \mathbb{E}_{x,y}\{\phi_1'(f(\rho,x),y)\nabla_1^T h(w,x)\} & 0 \end{bmatrix} \right\|_F \cdot \sqrt{trace(\nabla^2 r(\theta)\nabla^2 r(\theta)^T)} \\
&\leq \sqrt{dL_1^2 C_4^2\|\theta\|_2^2 + 2L_1^2 C_3^2} \cdot \sqrt{(d+1)D_3^2\|\theta\|_2^{2p} + 2(d+1)D_3D_4\|\theta\|_2^p + (d+1)D_4^2} \\
&\leq 2\sqrt{dL_1^2 C_4^2\|\theta\|_2^2 + 2L_1^2 C_3^2} \cdot \sqrt{(d+1)D_3^2\|\theta\|_2^{2p} + 2(d+1)D_3D_4\|\theta\|_2^p + (d+1)D_4^2} \\
&\leq (d+1)D_3^2\|\theta\|_2^{2p} + 2(d+1)D_3D_4\|\theta\|_2^p + dL_1^2 C_4^2\|\theta\|_2^2 + 2L_1^2 C_3^2 + (d+1)D_4^2 \ . \qquad \text{(B.28)}
\end{aligned}
$$

Hence, we obtain the following bound for $|\lambda_{min}(\nabla^2 U(\theta;\rho_t)|$,

$$
\begin{aligned}
&|\lambda_{min}(\nabla^2 U(\theta;\rho_t)| \\
&\leq \sqrt{2(d+1)D_3^2\|\theta\|_2^{2p} + 4(d+1)D_3D_4\|\theta\|_2^p + 2dL_1^2 C_4^2\|\theta\|_2^2 + 4L_1^2 C_3^2 + 2(d+1)D_4^2} \\
&\leq \sqrt{2(d+1)D_3^2}\|\theta\|_2^p + \sqrt{2dL_1^2 C_4^2}\|\theta\|_2 + 2\sqrt{(d+1)D_3D_4}\|\theta\|_2^{p/2} + \sqrt{4L_1^2 C_3^2 + 2(d+1)D_4^2} \ .
\end{aligned}
$$

Hence, we can choose

$$\Phi(\|\theta\|_2) = \sqrt{2(d+1)D_3^2}\|\theta\|_2^p + \sqrt{2dL_1^2C_4^2}\|\theta\|_2 + 2\sqrt{(d+1)D_3D_4}\|\theta\|_2^{p/2} + \sqrt{4L_1^2C_3^2 + 2(d+1)D_4^2} \ .$$

For $\mathcal{L}_t = \nabla \cdot \nabla - \nabla U(\theta; \rho_t) \cdot \nabla$, we now prove the following Lyapunov condition for the LSI, i.e., there is some nonnegative and locally Lipschitz function $V : \mathbb{R}^{d+1} \to \mathbb{R}^+ \cup \{0\}$ such that

$$\mathcal{L}_t V(\theta) + \|\nabla V(\theta)\|^2 \leq -c_1\|\theta\|_2^2\Phi(2\|\theta\|_2) + c_2 \ . \tag{B.29}$$

By setting $V(\theta) = \frac{\gamma}{2}\|\theta\|_2^2$ for some $\gamma > 0$ to be chosen later, we obtain

$$\mathcal{L}_t V(\theta) + \|\nabla V(\theta)\|^2 = \gamma(d+1) - \gamma\langle \nabla U(\theta; \rho_t), \theta\rangle + \gamma^2\|\theta\|_2^2 \ . \tag{B.30}$$

Then we have

$$\begin{aligned}
&\langle \nabla U(\theta; \rho_t), \theta\rangle \\
&= \left\langle \mathbb{E}_{x,y}\left\{ \phi_1'(f(\rho, x), y) \begin{bmatrix} u\nabla_1 h(w, x) \\ \\ h(w, x) \end{bmatrix} \right\} + \nabla r(\theta), \theta \right\rangle \\
&\geq \langle \nabla r(\theta), \theta\rangle - \left| \left\langle \mathbb{E}_{x,y}\left\{ \phi_1'(f(\rho, x), y) \begin{bmatrix} u\nabla_1 h(w, x) \\ \\ h(w, x) \end{bmatrix} \right\}, \theta \right\rangle \right| \\
&\geq m\|\theta\|_2^{2+p} - b - L_1\|\theta\|_2\sqrt{|h(w, x)|^2 + u^2\|\nabla_1 h(w, x)\|_2^2} \\
&\geq m\|\theta\|_2^{2+p} - b - L_1\|\theta\|_2|h(w, x)| - L_1\|\theta\|_2|u|\|\nabla_1 h(w, x)\|_2^2 \\
&\geq m\|\theta\|_2^{2+p} - b - L_1\|\theta\|_2(C_1\|w\| + C_2) - L_1C_3^2\|\theta\|_2^2 \\
&\geq m\|\theta\|_2^{2+p} - (L_1C_1 + L_1C_3^2)\|\theta\|_2^2 - L_1C_2\|\theta\|_2 - b \ .
\end{aligned} \tag{B.31}$$

Hence,

$$\begin{aligned}
&\mathcal{L}_t V(\theta) + \|\nabla V(\theta)\|^2 \\
&\leq \gamma(d+1) - \gamma m\|\theta\|_2^{2+p} + \gamma(\gamma + L_1C_1 + L_1C_3^2)\|\theta\|_2^2 + \gamma L_1C_2\|\theta\|_2 + \gamma b
\end{aligned} \tag{B.32}$$

Now, we consider to separate the whole space $\mathbb{R}^{d+1} = B(0, R) \cup (\mathbb{R}^{d+1} \setminus B(0, R))$ where

$$B(0, R) = \{\theta \in \mathbb{R}^{d+1} : \|\theta\|_2 \leq R\}$$

for some $R > 1$.

When $\theta \in B(0, R)$, we have

$$\begin{aligned}
&\mathcal{L}_t V(\theta) + \|\nabla V(\theta)\|^2 \\
&\leq \gamma(d+1) + \gamma(\gamma + L_1C_1 + L_1C_3^2)R^2 + \gamma L_1C_2R + \gamma b
\end{aligned} \tag{B.33}$$

and

$$-c_1\|\theta\|_2^2\Phi(2\|\theta\|_2) + c_2 \geq -c_1R^2\Phi(2R^2) + c_2 \ . \tag{B.34}$$

Hence, without loss of generality, we can set $c_1 = 1$ and

$$\begin{aligned}
c_2 &= c_2(R) = c_1 R^2 \Phi(2R) + \gamma(d+1) + \gamma(\gamma + L_1 C_1 + L_1 C_3^2) R^2 + \gamma L_1 C_2 R + \gamma b \\
&= (2R)^{p/2} \sqrt{2(d+1) D_3^2} + (2R) \sqrt{2 d L_1^2 C_4^2} + 2(2R)^{p/4} \sqrt{(d+1) D_3 D_4} \\
&+ R^2 \sqrt{4 L_1^2 C_3^2 + 2(d+1) D_4^2} + \gamma(d+1) + \gamma(\gamma + L_1 C_1 + L_1 C_3^2) R^2 + \gamma L_1 C_2 R + \gamma b \quad \text{(B.35)}
\end{aligned}$$

Then we have

$$\begin{aligned}
&\mathcal{L}_t V(\theta) + \|\nabla V(\theta)\|^2 \\
&\leq \gamma(d+1) + \gamma(\gamma + L_1 C_1 + L_1 C_3^2) R^2 + \gamma L_1 C_2 R + \gamma b \\
&= c_2(R) - c_1 R^2 \Phi(2R) \\
&\leq -c_1 \|\theta\|_2^2 \Phi(2\|\theta\|_2) + c_2 \ . \quad \text{(B.36)}
\end{aligned}$$

When $\theta \in B^c(0, R) = \mathbb{R}^{d+1} \setminus B(0, R)$,

$$\begin{aligned}
&\mathcal{L}_t V(\theta) + \|\nabla V(\theta)\|^2 \\
&\leq -\gamma m \|\theta\|_2^{2+p} + \gamma(\gamma + L_1 C_1 + L_1 C_3^2) \|\theta\|_2^2 + \gamma L_1 C_2 \|\theta\|_2 + \gamma(b + d + 1) \ . \quad \text{(B.37)}
\end{aligned}$$

Now, for $p > 1$, we can set any

$$\begin{aligned}
R^2 \geq \max \Bigg\{ & 1, \left( \frac{\sqrt{4 L_1^2 C_3^2 + 2(d+1) D_4^2} + \gamma(\gamma + L_1 C_1 + L_1 C_3^2)}{m} \right)^{\frac{2}{p}}, \\
& \left( \frac{\gamma L_1 C_2}{m} \right)^{\frac{2}{1+p}}, \left( \frac{\gamma(b+d+1)}{m} \right)^{\frac{2}{2+p}}, \left( \frac{2 L_1 C_4 \sqrt{d}}{m} \right)^{\frac{2}{p-1}}, \left( \frac{2\sqrt{(d+1) D_3 D_4}}{m} \right)^{\frac{4}{p}} \Bigg\}
\end{aligned}$$

and let

$$\gamma = \frac{1}{m} 2^{p/2} \sqrt{2(d+1)} D_3 + 5 \ . \quad \text{(B.38)}$$

Then we have the following bound for any $\theta \in B^c(0, R)$,

$$\begin{aligned}
&m \|\theta\|_2^{2+p} \\
&= m \|\theta\|_2^p \|\theta\|_2^2 \\
&\geq m R^p \|\theta\|_2^2 \\
&\geq \sqrt{4 L_1^2 C_3^2 + 2(d+1) D_4^2} \|\theta\|_2^2 + \gamma(\gamma + L_1 C_1 + L_1 C_3^2) \|\theta\|_2^2 \ . \quad \text{(B.39)}
\end{aligned}$$

Similarly, we also have the following bound for any $\theta \in B^c(0, R)$,

$$m \|\theta\|_2^{2+p} \geq 2\sqrt{d} L_1 C_4 \|\theta\|_2^3, \quad \text{(B.40)}$$

$$m \|\theta\|_2^{2+p} \geq 2\sqrt{(d+1) D_3 D_4} \|\theta\|_2^{p/2+2}, \quad \text{(B.41)}$$

$$m \|\theta\|_2^{2+p} \geq \gamma L_1 C_2 \|\theta\|_2, \quad \text{(B.42)}$$

$$m\|\theta\|_2^{2+p} \geq \gamma(b+d+1) \ . \tag{B.43}$$

Thus,

$$\begin{aligned}
&\mathcal{L}_t V(\theta) + \|\nabla V(\theta)\|^2 \\
&\leq -\gamma m\|\theta\|_2^{2+p} + \gamma(\gamma + L_1C_1 + L_1C_3^2)\|\theta\|_2^2 + \gamma L_1C_2\|\theta\|_2 + \gamma(b+d+1) \\
&= -2^{p/2}\sqrt{2(d+1)}D_3\|\theta\|_2^{p+2} - 5m\|\theta\|_2^{2+p} + \gamma(\gamma + L_1C_1 + L_1C_3^2)\|\theta\|_2^2 + \gamma L_1C_2\|\theta\|_2 + \gamma(b+d+1) \\
&= -2^{p/2}\sqrt{2(d+1)}D_3\|\theta\|_2^{p+2} - 2\sqrt{d}L_1C_4\|\theta\|_2^3 - 2\sqrt{(d+1)D_3D_4}\|\theta\|_2^{p/2+2} - \sqrt{4L_1^2C_3^2 + 2(d+1)D_4^2}\|\theta\|_2^2 \\
&= -\|\theta\|_2^2\Phi(2\|\theta\|_2) \ .
\end{aligned} \tag{B.44}$$

Similarly, for $p = 1$, we can set any

$$\begin{aligned}
R^2 \geq \max\Bigg\{ &1, \left(\frac{\sqrt{4L_1^2C_3^2 + 2(d+1)D_4^2} + \gamma(\gamma + L_1C_1 + L_1C_3^2)}{m}\right)^{\frac{2}{p}}, \\
&\left(\frac{\gamma L_1C_2}{m}\right)^{\frac{2}{1+p}}, \left(\frac{\gamma(b+d+1)}{m}\right)^{\frac{2}{2+p}}, \left(\frac{2\sqrt{(d+1)D_3D_4}}{m}\right)^{\frac{4}{p}}\Bigg\}
\end{aligned}$$

and let

$$\gamma = \frac{1}{m}2^{p/2}\left(\sqrt{2(d+1)}D_3 + 2L_1C_4\sqrt{d}\right) + 4 \ . \tag{B.45}$$

Then we can also get,

$$\begin{aligned}
&\mathcal{L}_t V(\theta) + \|\nabla V(\theta)\|^2 \\
&\leq -\gamma m\|\theta\|_2^{2+p} + \gamma(\gamma + L_1C_1 + L_1C_3^2)\|\theta\|_2^2 + \gamma L_1C_2\|\theta\|_2 + \gamma(b+d+1) \\
&= -2^{p/2}(\sqrt{2(d+1)}D_3 + 2L_1C_4\sqrt{d})\|\theta\|_2^{p+2} - 5m\|\theta\|_2^{2+p} \\
&+\gamma(\gamma + L_1C_1 + L_1C_3^2)\|\theta\|_2^2 + \gamma L_1C_2\|\theta\|_2 + \gamma(b+d+1) \\
&= -2^{p/2}(\sqrt{2(d+1)}D_3 + 2L_1C_4\sqrt{d})\|\theta\|_2^{p+2} - 2\sqrt{(d+1)D_3D_4}\|\theta\|_2^{p/2+2} - \sqrt{4L_1^2C_3^2 + 2(d+1)D_4^2}\|\theta\|_2^2 \\
&= -\|\theta\|_2^2\Phi(2\|\theta\|_2) \ .
\end{aligned} \tag{B.46}$$

In summary, for any $p \geq 1$, and $\theta \in \mathbb{R}^d$, if we set $V(\theta) = \frac{\gamma}{2}\|\theta\|_2^2$,we have

$$\mathcal{L}_t V(\theta) + \|\nabla V(\theta)\|^2 \leq -\|\theta\|_2^2\Phi(2\|\theta\|_2) + c_2 \ . \tag{B.47}$$

where

$$\begin{aligned}
c_2 = c_2(R) &:= c_1R^2\Phi(2R) + \gamma(d+1) + \gamma(\gamma + L_1C_1 + L_1C_3^2)R^2 + \gamma L_1C_2R + \gamma b \\
&= (2R)^{p/2}\sqrt{2(d+1)D_3^2} + (2R)\sqrt{2dL_1^2C_4^2} + 2(2R)^{p/4}\sqrt{(d+1)D_3D_4} \\
&+R^2\sqrt{4L_1^2C_3^2 + 2(d+1)D_4^2} + \gamma(d+1) + \gamma(\gamma + L_1C_1 + L_1C_3^2)R^2 + \gamma L_1C_2R + \gamma b
\end{aligned}$$

and when $p = 1$, we can set any

$$R^2 \geq \max\left\{1, \left(\frac{\sqrt{4L_1^2C_3^2 + 2(d+1)D_4^2} + \gamma(\gamma + L_1C_1 + L_1C_3^2)}{m}\right)^{\frac{2}{p}}, \right.$$
$$\left.\left(\frac{\gamma L_1 C_2}{m}\right)^{\frac{2}{1+p}}, \left(\frac{\gamma(b+d+1)}{m}\right)^{\frac{2}{2+p}}, \left(\frac{2\sqrt{(d+1)D_3D_4}}{m}\right)^{\frac{4}{p}}\right\}$$

and let

$$\gamma = \frac{1}{m}2^{p/2}\left(\sqrt{2(d+1)}D_3 + 2L_1C_4\sqrt{d}\right) + 4\ ; \tag{B.48}$$

when $p > 1$, we can set any

$$R^2 \geq \max\left\{1, \left(\frac{\sqrt{4L_1^2C_3^2 + 2(d+1)D_4^2} + \gamma(\gamma + L_1C_1 + L_1C_3^2)}{m}\right)^{\frac{2}{p}}, \right.$$
$$\left.\left(\frac{\gamma L_1 C_2}{m}\right)^{\frac{2}{1+p}}, \left(\frac{\gamma(b+d+1)}{m}\right)^{\frac{2}{2+p}}, \left(\frac{2L_1C_4\sqrt{d}}{m}\right)^{\frac{2}{p-1}}, \left(\frac{2\sqrt{(d+1)D_3D_4}}{m}\right)^{\frac{4}{p}}\right\}$$

and let

$$\gamma = \frac{1}{m}2^{p/2}\sqrt{2(d+1)}D_3 + 5\ . \tag{B.49}$$

□

Now, for $t \geq 0$, we are ready to prove Theorem 2, obtaining the following uniform bound for $\nu_{\rho_t}$.

**Theorem 4.** *We have the following uniform bound on the Logarithmic Sobolev constant $\nu_{\rho_t}$.*

$$\sup_{t\geq 0}\nu_{\tilde{\rho}_t} \leq 2\sqrt{\frac{2}{c_1\lambda^{-1}\Phi(0)}} + \frac{1}{c_1}$$
$$+\left[2\sqrt{\frac{2}{c_1\lambda^{-1}\Phi(0)}}c_2 + \frac{2c_2}{c_1} + 2\right]\inf_{r>\sqrt{\frac{c_2}{c_1\lambda^{-1}\Phi(0)}}}\frac{1 + Cc_1r^4\lambda^{-1}\Phi(0)e^{2poly(r)+\gamma r^2/2}}{c_1r^2\lambda^{-1}\Phi(0) - c_2} := \nu$$

*where $poly(r) = C_7r^k + L_1C_1r^2 + L_1C_2r + C_8$.*

*Proof.* For any $t \geq 0$, we have

$$\begin{aligned}
&|U(\rho_t;\theta)| \\
&= |\mathbb{E}_{x,y}\phi_1'(f(\rho,x),y)\tilde{h}(\theta,x) + r(\theta)| \\
&\leq L_1\|\theta\|_2(C_1\|\theta\|_2 + C_2) + C_7\|\theta\|_2^k + C_8 \\
&= C_7\|\theta\|_2^k + L_1C_1\|\theta\|_2^2 + L_1C_2\|\theta\|_2 + C_8
\end{aligned} \tag{B.50}$$

Hence, we obtain

$$\sup_{\|\theta\|_2 \le r} |U(\rho_t;\theta)| \le C_7 r^k + L_1 C_1 r^2 + L_1 C_2 r + C_8 := poly(r) \tag{B.51}$$

and

$$Osc_r(U_t) \le 2 \sup_{\|\theta\|_2 \le r} |U(\rho_t;\theta)| \le 2(C_7 r^k + L_1 C_1 r^2 + L_1 C_2 r + C_8) = 2poly(r)\ . \tag{B.52}$$

□

The following property of the operator $\mathcal{L}_t$ is useful in establishing the DLSI in our analysis. It appears in many literatures about LSI, such as Cattiaux et al. (2009) and Cattiaux et al. (2010), but we include the proof here for completeness.

**Lemma 3.** *Let $V$ be a nonnegative and locally Lipschitz function such that $\mathcal{L}_t V + \|\nabla V\|_2^2 \le -\phi$, where $\phi$ is lower bounded. Then*

$$\int \phi g^2 d\mu_t \le \mathcal{E}_t(g), \quad g \in \mathbb{D}_2(\mathcal{L}_t). \tag{B.53}$$

*Proof.* Since $\phi \wedge N$ also satisfies the condition and can approach $\phi$ by letting $N \to \infty$, we can assume that $\phi$ is bounded. Furthermore, since $\mathbb{D}_2(\mathcal{L}_t)$ is a dense subset of $L^2(\mu_t)$, we can approach any $g \in \mathbb{D}_2(\mathcal{L}_t)$ by $(g_n) \subset C_0^\infty(\mathbb{R}^{d+1}) : \|g_n - g\|^2_{L^2(\mu_t)} + \mathcal{E}_t(g_n - g) \to 0$. Hence it is sufficient to prove for $g \in C_0^\infty(\mathbb{R}^{d+1})$. By the self adjoint property, we have

$$\langle -\mathcal{L}_t V, g^2 \rangle_{L^2(\mu_t)} = \langle V, -\mathcal{L}_t g^2 \rangle_{L^2(\mu_t)} = \langle \nabla V, \nabla(g^2) \rangle_{L^2(\mu_t)} = \int_{\mathbb{R}^{d+1}} \nabla V \cdot \nabla(g^2) d\mu_t\ . \tag{B.54}$$

Therefore, we have

$$\begin{aligned}
&\int_{\mathbb{R}^{d+1}} \phi g^2 d\mu_t \\
&\le \int_{\mathbb{R}^{d+1}} (-\mathcal{L}_t V - \|\nabla V\|_2^2) g^2 d\mu_t \\
&= \int_{\mathbb{R}^{d+1}} (\nabla V \cdot \nabla(g^2) - \|\nabla V\|_2^2 g^2) d\mu_t \\
&= \int_{\mathbb{R}^{d+1}} (2g \nabla V \cdot \nabla g - \|\nabla V\|_2^2 g^2) d\mu_t \\
&\le \int_{\mathbb{R}^{d+1}} \|\nabla g\|_2^2 d\mu_t = \mathcal{E}_t(g)
\end{aligned} \tag{B.55}$$

where in the last inequality, we use $2g\nabla V \cdot \nabla g \le \|\nabla V\|_2^2 g^2 + \|\nabla g\|_2^2$. □

# C Experimental Results

We conduct experiments to demonstrate the role of different regularizers to the convergence.

In the experiment, we consider $d = 2$ and the data is generated from a two-layer teacher neural network with two hidden neurons $w_1 = [1, 2]^T, u_1 = 1.1, w_2 = [-3, 1]^T, u_2 = -3.2$. We choose smoothed-ReLU as activations according to our theoretic results. We consider an overparametrized student network with 20 hidden neurons. The weights are initialized by standard Gaussian. Initialized weights are shown in Figure 1.

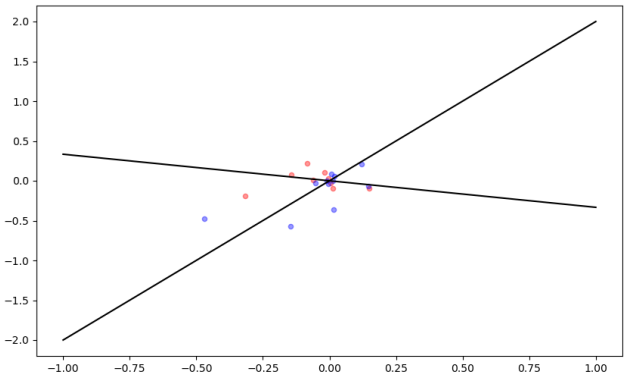


Figure 1: Initialized neurons can be visualized in a two-dimensional plane by $uw \in \mathbb{R}^2$. Positive wighted neurons ($u \geq 0$) are in red and negative weighted neurons ($u < 0$) are in blue. Two black lines are the two hidden neurons of the teacher network.

We choose the regularized square loss, with $\beta = 1.0$ and $\lambda = 1.0$. The initial learning rate is 0.0001 and decays every 100 epochs. We trained the student network for 200 epochs and we generate 200 data points from the teacher network to train the student network in each epoch.

We train the two student networks with the same initialized weights but different regularizers: $r_1(\theta) = \frac{\beta}{2}\|\theta\|_2^2$ for one network and $r_2(\theta) = \frac{\beta}{3}\|\theta\|_2^3$ for another network.

Figure 2 shows the evolution of unregularized square loss and Figure 3 shows the regularized square loss of two student networks during the training. The learned weights for two networks are shown in Figure 4 and Figure 5 respectively.

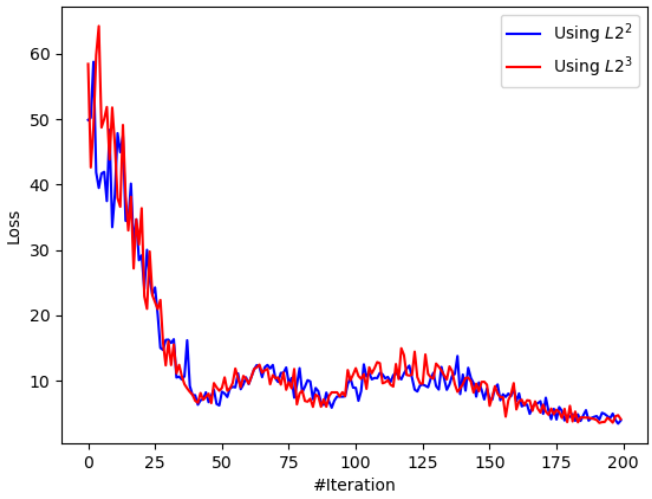


Figure 2: Square Loss

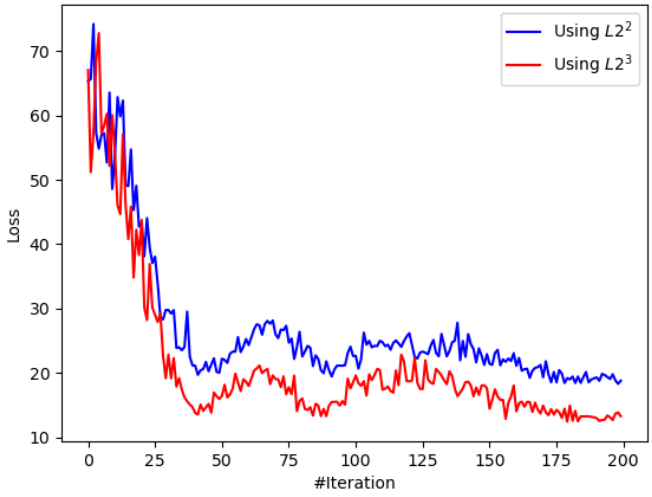


Figure 3: Regularized Square Loss

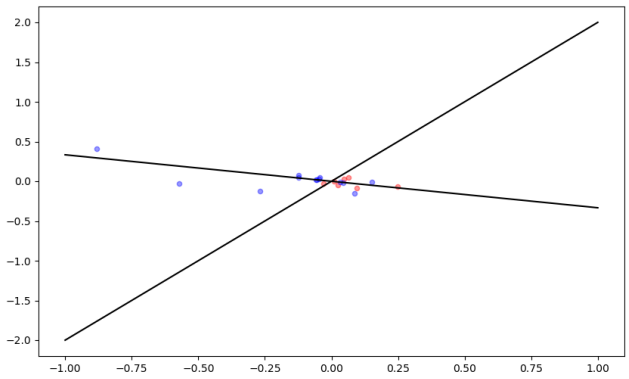


Figure 4: Learned Weights for $r_1(\theta) = \frac{\beta}{2}\|\theta\|_2^2$

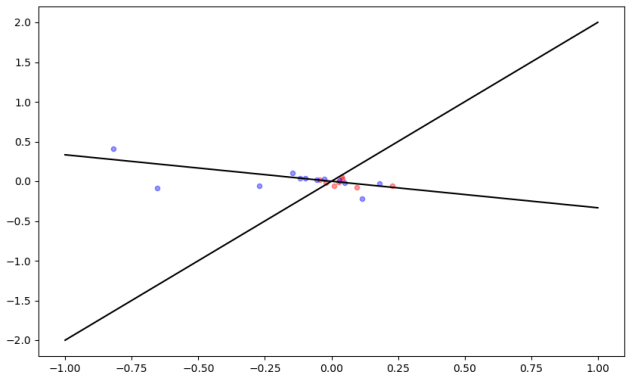


Figure 5: Learned Weights for $r_2(\theta) = \frac{\beta}{3}\|\theta\|_2^3$

From the experiments, we know that the cubic regularizer $r_2(\theta)$ has similar convergence to $r_1(\theta)$ for both the regularized objective as shown in Figure 3 and the unregularized objective as shown in Figure 2. Furthermore, the results demonstrate that both regularizers can learn weights and positions efficiently in the mean-field regime while the NTK regime suffers from lazy training and can only learn the weight $u$ of hidden neurons. The experiments show that the cubic regularizer achieves comparable performance to the original weight decay, with additional theoretic linear convergence guarantees in the many-particle and continuous time limit.

# D Further Discussions

## D.1 Intuitions behind Theorem 1 and Theorem 2

As shown in our analysis, establishing the uniform LSI for $\tilde{\rho}_t \propto \exp(-\lambda^{-1}U(\theta, \rho_t))$ along the trajectory is the key technical contribution to obtain the linear convergence of two-layer NN trained

by noisy SGD in the mean-field regime. This is purely a math problem and we need to adopt the literature of LSI, where a general principle is to control the curvature lower bound of $U(\theta, \rho_t)$ to obtain the LSI.

The most fundamental idea to establish the LSI is by the Baker-Emery criteria, when the potential $U(\theta, \rho_t))$ is a strongly convex function of $\theta$ (i.e., with positive curvature lower bound). However, for neural networks, the convexity of potential $U(\theta, \rho_t))$ is hard to satisfy. One can generalize the Bakery-Emery criteria to the nonconvex potentials by the Holley-Stroock perturbation lemma: if the nonconvex potential is a bounded perturbation of some strongly convex potential, then it also satisfies the LSI.

This is exactly what [Nitanda et al. '22] & [Chizat '22] adopts to establish LSI: they assume that the neural networks is a bounded perturbation of the $\|\theta\|_2^2$ regularizer. From a technical perspective, treating the neural network as the perturbation of the regularizer seems not satisfactory. In this sense, the convergence analysis is a result from the regularizer. The bounded assumption is not satisfied by the vanilla two-layer network unless we fix the parameter of the second layer.

Hence, Theorem 1 in this paper can be considered as a nontrivial extension to the result of [Nitanda et al. '22] & [Chizat '22], where we consider the regularizer with growth at least $\|\theta\|_2^4$. In such sense, there is more "room" for nonconvex perturbation and the "room" can fit the vanilla two-layer neural networks. As shown in our derivations, we see that the Hessian of the nonconvex potential contributed by neural networks grows with $\mathcal{O}(-\|\theta\|_2 I)$ whereas the Hessian of the regularizer $\|\theta\|_2^p$ grow with $\mathcal{O}(\|\theta\|_2^{p-2} I)$. Hence as long as $p-2>1$, that is $p \geq 4$, we can always ensure the Hessian of the nonconvex potential be positive outside a ball in $\mathbb{R}^{d+1}$.

Finally in Theorem 2, we aim to find a tightest possible "room" that can fit vanilla two-layer neural networks. Therefore, we adopt the recent state-of-art results in the LSI literature to estimate the Log-Sobolev constant for vanilla two-layer neural networks. In particular, Cattiaux et al. (2010) obtains that when the lower bound on the Hessian of the potential diverges to $-\infty$, but can be controlled by some nondecreasing function $\Phi(\theta)$, then the LSI can also be obtained under the Lyapunov conditions in (33).

For the behavior of the right-hand side of the inequality in two theorems, one could observe that there exists an optimal $\lambda$ for the entropy regularizer that produce the sharpest bound. One could use this optimal $\lambda$ in model selection, where we can choose the most suitable noise level in noisy SGD to accelerate the convergence. Another take-home message is that smoother loss and activations guarantee better convergence, which is also desirable intuitively.

### D.2 Is $p \geq 3$ for $r(\theta) = \beta\|\theta\|_2^p$ necessary?

After some careful analysis on the potential contributed by two-layer neural networks and Lyapunov conditions in equation (33), we can easily found that the best choice of $\Phi(\theta)$ is of order $\|\theta\|_2$ and the weakest regularization is of order $\|\theta\|_2^3$ to establish the Lyapunov condition in (33).

If we want to relax the restrictions on $p \geq 3$, we need to develop new mathematical tools to establish LSI better than Cattiaux et al. (2010), which should be the job for mathematicians.