

LEARNING AND GENERALIZATION IN UNIVARIATE OVERPARAMETERIZED NORMALIZING FLOWS

Anonymous authors

Paper under double-blind review

ABSTRACT

In supervised learning, it is known that overparameterized neural networks with one hidden layer provably and efficiently learn and generalize, when trained using Stochastic Gradient Descent (SGD). In contrast, the benefit of overparameterization in unsupervised learning is not well understood. Normalizing flows (NFs) learn to map complex real-world distributions into simple base distributions, and constitute an important class of models in unsupervised learning for sampling and density estimation. In this paper, we theoretically and empirically analyze these models when the underlying neural network is one hidden layer overparametrized network. On the one hand we provide evidence that for a class of NFs, overparametrization hurts training. On the other, we prove that another class of NFs, with similar underlying networks can efficiently learn any reasonable data distribution under minimal assumptions. We extend theoretical ideas on learning and generalization from overparameterized neural networks in supervised learning to overparameterized normalizing flows in unsupervised learning. We also provide experimental validation to support our theoretical analysis in practice.

1 INTRODUCTION

Neural network models trained using simple first-order iterative algorithms have been very effective in both supervised and unsupervised learning. Theoretical reasoning of this phenomenon requires one to consider simple but quintessential formulations, where this can be demonstrated by a mathematical proof, along with experimental evidence for the underlying intuition. First, the minimization of training loss is typically a non-smooth and non-convex optimization over the parameters of neural networks, so it is surprising that neural networks can be trained efficiently by first-order iterative algorithms. Second, even large neural networks whose number parameters are more than the size of training data often generalize well with small loss on the unseen test data, instead of overfitting the seen training data. Recent work in supervised learning attempts to provide theoretical justification for why overparameterized neural networks can train and generalize efficiently in the above sense.

In supervised learning, the empirical risk minimization with quadratic loss is a non-convex optimization problem even for a fully connected neural network with one hidden layer of neurons with ReLU activations. Around 2018, it was realized that when the hidden layer size is large compared to the dataset size or compared to some measure of complexity of the data, one can provably show efficient training and generalization for these networks, e.g. Jacot et al. (2018); Li & Liang (2018); Du et al. (2018); Allen-Zhu et al. (2019); Arora et al. (2019). Of these, Allen-Zhu et al. (2019) is directly relevant to our paper and will be discussed later.

The role of overparameterization, and provable training and generalization guarantees for neural networks are less well understood in unsupervised learning. Generative models or learning a data distribution from given samples is an important problem in unsupervised learning. Popular generative models based on neural networks include Generative Adversarial Networks (GANs) (e.g., Goodfellow et al. (2014)), Variational AutoEncoders (VAEs) (e.g., Kingma & Welling (2014)), and Normalizing Flows (e.g., Rezende & Mohamed (2015)). GANs and VAEs have shown impressive capability to generate samples of photo-realistic images but they cannot give probability density estimates for new data points. Training of GANs and VAEs has various additional challenges such as mode collapse, posterior collapse, vanishing gradients, training instability etc. as shown in e.g. Bowman et al. (2016); Salimans et al. (2016); Arora et al. (2018); Lucic et al. (2018). Recent work has studied learning

and generalization in GANs. Lei et al. (2020) show that when the generator is a one-layer tanh, sigmoid or leaky ReLU network, Wasserstein GAN training with stochastic gradient descent-ascent converges to a global solution with polynomial time and sample complexity. Using the method of moments and a learning algorithm motivated by tensor decomposition, Li & Dou (2020) show that GANs can efficiently learn a large class of distributions including those generated by two-layer networks. Recent work has also looked at provably efficient unsupervised learning in other neural generative models beside GANs. Nguyen et al. (2019b) show that two-layer autoencoders with ReLU or threshold activations can be trained with normalized gradient descent over the reconstruction loss to provably learn the parameters of any generative bilinear model (e.g., mixture of Gaussians, sparse coding model). Nguyen et al. (2019a) extend the work of Du et al. (2018) on supervised learning mentioned earlier to study weakly-trained (i.e., only encoder is trained) and jointly-trained (i.e., both encoder and decoder are trained) two-layer autoencoders, and show joint training requires less overparameterization and converges to a global optimum. The effect of overparameterization in unsupervised learning has also been of recent interest. Buhai et al. (2020) do an empirical study to show that across a variety of latent variable models (e.g., single-layer models with neural parameterization) and training algorithms, overparameterization can significantly increase the number of recovered ground truth latent variables. Radhakrishnan et al. (2020) show that overparameterized autoencoders and sequence encoders essentially implement associative memory by storing training samples as attractors in a dynamical system.

In contrast to the generative models such as GANs and VAEs, when normalizing flows learn distributions, they can do both sampling and density estimation, leading to wide-ranging applications as mentioned in the surveys by Kobzyev et al. (2020) and Papamakarios et al. (2019). Previous work has studied different variants of normalizing flows such as planar and radial flows in Rezende & Mohamed (2015), Sylvester flow in van den Berg et al. (2018), Householder flow in Tomczak & Welling (2016), masked autoregressive flow in Papamakarios et al. (2017). Most variants of normalizing flows are specific to certain applications, and the expressive power (i.e., which base and data distributions they can map between) and complexity of normalizing flow models has been studied recently, e.g. Kong & Chaudhuri (2020) and Teshima et al. (2020). Invertible transformations defined by monotonic neural networks can be combined into autoregressive flows that are universal density approximators of continuous probability distributions; see Masked Autoregressive Flows (MAF) Papamakarios et al. (2017), UNMM-MAF by Wehenkel & Louppe (2019), Neural Autoregressive Flows (NAF) by Huang et al. (2018), Block Neural Autoregressive Flow (B-NAF) by Cao et al. (2019). Unconstrained Monotonic Neural Network (UMNN) models proposed by Wehenkel & Louppe (2019) are particularly relevant to the technical part of our paper. Theoretical understanding of learning and generalization in normalizing flows (more generally, generative models and unsupervised learning) is a natural and important open question, and our main technical contribution is to extend known techniques from supervised learning to make progress towards answering this question. In this paper, we study learning and generalization in the case of univariate overparameterized normalizing flows. Restriction to the univariate case is technically non-trivial and interesting in its own right: univariate ReLU networks have been studied in recent supervised learning literature (e.g., Savarese et al. (2019), Williams et al. (2019), Sahs et al. (2020) and Daubechies et al. (2019)). Multidimensional flows are qualitatively more complex and our 1D analysis sheds some light on them (see Sec. 4). Before stating our contributions, we briefly introduce normalizing flows; details appear in Sec. 2.

Normalizing Flows. We work with one-dimensional probability distributions with continuous density. The general idea behind normalizing flows (NFs), restricted to 1D can be summarized as follows: Let $X \in \mathbb{R}$ be a random variable denoting the data distribution. We also fix a *base* distribution with associated random variable Z which is typically standard Gaussian, though in this paper we will work with the exponential distribution as well. Given i.i.d. samples of X , the goal is to learn a continuous strictly monotone increasing map $f_X : \mathbb{R} \rightarrow \mathbb{R}$ that transports the distribution of X to the distribution of Z : in other words, the distribution of $f_X^{-1}(Z)$ is that of X . The learning of f_X is done by representing it by a neural network, and setting up an appropriate loss function.

The monotonicity requirement on f which makes f invertible, while not essential, greatly simplifies the problem and is present in all the works we are aware of. It is not clear how to set up a tractable optimization problem without this requirement. Since the function represented by standard neural networks are not necessarily monotone, the design of the neural net is altered to make it monotone. For our 1D situation, one-hidden layer networks are of the form $N(x) = \sum_{i=1}^m a_i \sigma(w_i x + b_i)$, where m is the size of the hidden layer and the a_i, w_i, b_i are the parameters of the network.

We will assume that the activation functions used are monotone. Here we distinguish between two such alterations: (1) Changing the parametrization of the neural network. This can be done in multiple ways: instead of a_i, w_i we use a_i^2, w_i^2 (or other functions, such as the exponential function, of a_i, w_i that take on only positive values) (Huang et al., 2018; Cao et al., 2019). This approach appears to be the most popular. In this paper, we also suggest another related alteration: we simply restrict the parameters a_i, w_i to be positive. This is achieved by enforcing this constraint during training. (2) Instead of using $N(x)$ for $f(x)$ we use $\phi(N(x))$ for $f'(x) = \frac{df}{dx}$, where $\phi: \mathbb{R} \rightarrow \mathbb{R}^+$ takes on only positive values. Positivity of f' implies monotonicity of f . Note that no restrictions on the parameters are required; however, because we parametrize f' , the function f needs to be reconstructed using numerical quadrature. This approach is used by Wehenkel & Louppe (2019).

We will refer to the models in the first class as *constrained normalizing flows* (CNFs) and those in the second class as *unconstrained normalizing flows* (UNFs).

Our Contributions. In this paper we study both constrained and unconstrained univariate NFs theoretically as well as empirically.

The existing analyses for overparametrized neural networks in the supervised setting work with a linear approximation of the neural network, termed *pseudo network* in Allen-Zhu et al. (2019). They show that (1) there is a pseudo network with weights close to the initial ones approximating the target function, (2) the loss surfaces of the neural network and the pseudo network are close and moreover the latter is convex for convex loss functions. This allows for a proof of the convergence of the training of neural network to global optima.

One can try to adapt the approach of using a linear approximation of the neural network to analyze training of NFs. However, one immediately encounters some new roadblocks: the loss surface of the pseudo networks is non-convex in both CNFs and UNFs.

In both cases, we identify novel variations that make the optimization problem for associated pseudo network convex: For CNFs, instead of using a_i^2, w_i^2 as parameters, we simply impose the constraints $a_i \geq \epsilon$ and $w_i \geq \epsilon$ for some small constant ϵ . The optimization algorithm now is projected SGD, which in this case incurs essentially no extra cost over SGD due to the simplicity of the positivity constraints. Apart from making the optimization problem convex, in experiments this variation slightly *improves* the training of NFs compared to the reparametrization approaches, and may be useful in practical settings.

Similarly, for UNFs we identify two changes from the model of Wehenkel & Louppe (2019) that make the associated optimization problem convex, while still retaining empirical effectiveness: (1) Instead of Clenshaw–Curtis quadrature employed in Wehenkel & Louppe (2019) which uses positive and negative coefficients, we use the simple rectangle quadrature which uses only positive coefficients. This change makes the model somewhat slow (it uses twice as many samples and time to get similar performance on the examples we tried). (2) Instead of the standard Gaussian distribution as the base distribution, we use the exponential distribution. In experiments, this does not cause much change.

Our results point to a dichotomy between these two classes of NFs: our variant of UNFs can be theoretically analyzed when the networks are overparametrized to prove that the UNF indeed learns the data distribution. To our knowledge, this is the first “end-to-end” analysis of an NF model, and of a neural generative model using gradient-based algorithms used in practice. This proof, while following the high-level scheme of Allen-Zhu et al. (2019) proof, has a number of differences, conceptual as well as technical, due to different settings. E.g., our loss function involves a function and its integral estimated by quadrature.

On the other hand, for CNFs, our empirical and theoretical findings provide evidence that overparametrization makes training slower to the extent that models of similar size which learn the data distribution well for UNFs, fail to do so for CNFs. We also analyze CNFs theoretically in the overparametrized setting and point to potential sources of the difficulty. The case of moderate-sized networks, where training and generalization does take place empirically, is likely to be difficult to analyze theoretically as presently this setting is open for the simpler supervised learning case. We hope that our results will pave the way for further progress. We make some remarks on the multidimensional case in Sec. 4. In summary, our contributions are:

- To our knowledge, first efficient training and generalization proof for NFs (in 1D).

- Identification of architectural variants of UNFs that admit analysis via overparametrization.
- Identification of “barriers” to the analysis of CNFs.

A brief outline of our paper is as follows. Section 2 contains preliminaries and an overview of our results about constrained and unconstrained normalizing flows. Appendix A shows the existence of a pseudo network whose loss closely approximates the loss of the target function. Appendix B shows coupling or closeness of their gradients over random initialization. Appendices C and D contain complete proofs of our optimization and generalization results, respectively. Section 3 and Appendix E contain our empirical studies towards validating our theoretical results.

2 PRELIMINARIES AND OVERVIEW OF RESULTS

We confine our discussion to the 1D case which is the focus of the present paper. The goal of NF is to learn a probability distribution given via i.i.d. samples data. We will work with distributions whose densities have finite support, and assumed to be $[-1, 1]$, without loss of generality. Let X be the random variable corresponding to the data distribution we want to learn. We denote the probability density (we often just say density) of X at $u \in \mathbb{R}$ by $p_X(u)$. Let Z be a random variable with either standard Gaussian or the exponential distribution with $\lambda = 1$ (which we call standard exponential). Recall that the density of the standard exponential distribution at $u \in \mathbb{R}$ is given by e^{-u} for $u \geq 0$ and 0 for $u < 0$.

Let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a strictly increasing continuous function. Thus, f is invertible. We use $f'(x) = \frac{df}{dx}$ to denote the derivative. Let $p_{f,Z}(\cdot)$ be the density of the random variable $f^{-1}(Z)$. Let $x = f^{-1}(z)$, for $z \in \mathbb{R}$. Then by the standard change of density formula using the monotonicity of f gives

$$p_{f,Z}(x) = p_Z(z)f'(x). \quad (1)$$

We would like to choose f so that $p_{f,Z} = p_X$, the true data density. It is known that such an f always exists and is unique; see e.g. Chapter 2 of Santambrogio (2015). We will refer to the distribution of Z as the *base* distribution. Note that if we can find f , then we can generate samples of X using $f^{-1}(Z)$ since generating the samples of Z is easy. Similarly, we can evaluate $p_X(x) = p_Z(f^{-1}(z))f'(x)$ using (1). To find f from the data, we set up the maximum log-likelihood objective:

$$\max_f \frac{1}{n} \sum_{i=1}^n \log p_{f,Z}(x_i) = \max_f \frac{1}{n} \left[\sum_{i=1}^n \log p_Z(f(x_i)) + \sum_{i=1}^n \log f'(x_i) \right], \quad (2)$$

where $S = \{x_1, \dots, x_n\} \subset \mathbb{R}$ contains i.i.d. samples of X , and the maximum is over continuous strictly increasing functions. When Z is standard exponential, the optimization problem (2) becomes

$$\min_f L(f, S), \quad \text{where } L(f, S) = \frac{1}{n} \sum_{x \in S} L(f, x) \text{ and } L(f, x) = f(x) - \log f'(x). \quad (3)$$

A similar expression, with $f(x)^2/2$ replacing $f(x)$, holds for the standard Gaussian. We denote the loss for standard Gaussian as $L_G(f, x)$.

Informally, one would expect that as $n \rightarrow \infty$, for the optimum f in the above optimization problems $p_{f,Z} \rightarrow p_X$. To make the above optimization problem tractable, instead of f we use a neural network N . We consider one-hidden layer neural networks with the following basic form which will then be modified according to whether we are constraining the parameters or the output.

$$N(x) = \sum_{r=1}^m a_{r0} \rho((w_{r0} + w_r)x + (b_r + b_{r0})). \quad (4)$$

Here m is the size of the hidden layer, $\rho : \mathbb{R} \rightarrow \mathbb{R}$ is a monotonically increasing activation function, the weights a_{r0}, w_{r0}, b_{r0} are the initial weights chosen at random according to some distribution, and w_r, b_r are offsets from the initial weights. We will only train the w_r, b_r and the a_{r0} will remain frozen to their initial values.

Let $\theta = (W, B) \in \mathbb{R}^{2m}$ denote the parameters $W = (w_1, w_2, \dots, w_m) \in \mathbb{R}^m$ and $B = (b_1, b_2, \dots, b_m) \in \mathbb{R}^m$ of the neural network. We use Stochastic Gradient Descent (SGD) to update the parameters of neural networks. Denote by $\theta^t = (W^t, B^t)$ with $W^t = (w_1^t, w_2^t, \dots, w_m^t)$ and

$B^t = (b_1^t, b_2^t, \dots, b_m^t)$ the parameters at time step $t = 1, 2, \dots$, and the corresponding network by $N_t(x)$. The SGD updates are given by $\theta^{t+1} = \theta^t - \eta \nabla_{\theta} L(N_t, x^t)$ where $\eta > 0$ is learning rate, and $L(N_t, x^t)$ is a loss function, and $x^t \in S$ is chosen uniformly randomly at each time step. For supervised learning where we are given labeled data $\{(x_1, y_1), \dots, (x_n, y_n)\}$, one often works with the mean square loss $L(N_t) = \frac{1}{n} \sum_{i=1}^n L(N_t, x_i)$ with $L(N_t, x_i) = (N_t(x_i) - y_i)^2$.

We now very briefly outline the proof technique of Allen-Zhu et al. (2019) for analyzing training and generalization for one-hidden layer neural networks for supervised learning. (While they work in a general agnostic learning setting, for simplicity, we restrict the discussion to the realizable setting.) In their setting, the data $x \in \mathbb{R}^d$ is generated by some distribution D and the labels $y = h(x)$ are generated by some unknown function $h: \mathbb{R}^d \rightarrow \mathbb{R}$. The function h is assumed to have small ‘‘complexity’’ C_h which in this case measures the required size of neural network with smooth activations to approximate h .

The problem of optimizing the square loss is non-convex even for one-hidden layer networks. Allen-Zhu et al. (2019) instead work with *pseudo network*, $P(x)$ which is the linear approximation of $N(x)$ given by the first-order Taylor expansion of the activation:

$$P(x) = \sum_{r=1}^m a_{r0} (\sigma(w_{r0}x + b_{r0}) + \sigma'(w_{r0}x + b_{r0}) (w_r x + b_r)). \quad (5)$$

Similarly to N_t we can also define P_t with parameters θ^t . They observe that when the network is highly overparameterized, i.e. the network size m is sufficiently large compared to C_h , and the learning rate is small, i.e. $\eta = O(1/m)$, SGD iterates when applied to $L(N_t)$ and $L(P_t)$ remain close throughout. Moreover, the problem of optimizing $L(P)$ is a convex problem in θ and thus can be analyzed with existing methods. They also show an *approximation* theorem stating that with high probability there are neural network parameters θ^* close to the initial parameters θ^0 such that the pseudo network with parameters θ^* is close to the target function. This together with the analysis of SGD shows that the pseudo network, and hence the neural network too, achieves small training loss. Then by a Rademacher complexity argument they show that the neural network after $T = O(C_h/\epsilon^2)$ time steps has population loss within ϵ of the optimal loss, thus obtaining a generalization result.

We will now describe how to obtain neural networks representing monotonically increasing functions using the two different methods mentioned earlier, namely CNFs and UNFs.

2.1 CONSTRAINED NORMALIZING FLOW

Note that if we have $a_{r0} \geq 0$, $w_{r0} + w_r \geq 0$ for all r , then the function represented by the neural network is monotonically increasing. We can ensure this positivity constraint by replacing a_{r0} and $w_{r0} + w_r$ by their functions that take on only positive values. For example, the function $x \mapsto x^2$ would give us the neural network $N(x) = \sum_{r=1}^m a_{r0}^2 \rho((w_{r0} + w_r)^2 x + b_{r0} + b_r)$. Note that a_{r0} , $w_{r0} + w_r$ and $b_{r0} + b_r$ have no constraints, and so this network can be trained using standard gradient-based algorithms. But first we need to specify the (monotone) activation ρ . Let $\sigma(x) = x \mathbb{I}[x \geq 0]$ denote the ReLU activation. If we choose $\rho = \sigma$, then note that in (3) we have

$$\log f'(x) = \log \frac{\partial N(x)}{\partial x} = \log \left(\sum_{r=1}^m a_{r0}^2 (w_{r0} + w_r)^2 \mathbb{I}[(w_{r0} + w_r)^2 x + b_{r0} + b_r \geq 0] \right).$$

This is a discontinuous function in x as well as in w_r and b_r . Gradient-based optimization algorithms are not applicable to problems with discontinuous objectives, and indeed this is reflected in experimental failure of such models in learning the distribution. By the same argument, any activation that has a discontinuous derivative is not admissible. Activations which have continuous derivative but are not convex (e.g. ELU(x) given by $e^x - 1$ for $x < 0$ and x for $x \geq 0$) also cannot be used because then $N(x)$ is also a convex function of x , which need not be the case for the optimal f . The oft-used activation tanh does not suffer from either of these defects. Pseudo network is now given by

$$P(x) = \sum_{r=1}^m a_{r0}^2 (\tanh(w_{r0}^2 x + b_{r0}) + \tanh'(w_{r0}^2 x + b_{r0}) ((w_r^2 + 2w_{r0}w_r) x + b_r)).$$

Note that $P(x)$ is not linear in the parameters θ . Hence, it is not obvious that the loss function for pseudo network will remain convex in parameters and non-convexity can be confirmed in experiments. A similar situation arises for exponential parameterization instead of square.

To overcome the non-convexity issue, we propose another formulation for constrained normalizing flows. Here we retain the form of the neural network as in (4), but ensure the constraints $a_{r0} \geq 0$ and $w_{r0} \geq 0$ by the choice of the initialization distribution and $w_{r0} + w_r \geq 0$ by using projected gradient descent for optimization.

$$N(x) = \sum_{r=1}^m a_{r0} \tanh((w_{r0} + w_r)x + (b_r + b_{r0})), \text{ with constraints } w_{r0} + w_r \geq \epsilon, \text{ for all } r.$$

Here, $\epsilon > 0$ is a small constant ensuring strict monotonicity of $N(x)$. Note that constraints in the formulation are simple and easy to use in practice. The pseudo network in this formulation will be

$$P(x) = \sum_{r=1}^m a_{r0} (\tanh(w_{r0}x + b_{r0}) + \tanh'(w_{r0}x + b_{r0})(w_r x + b_r)),$$

with constraints $w_{r0} + w_r \geq \epsilon$, for all r . $P(x)$ is linear in θ , therefore the objective function is also convex in θ . Note that $P(x)$ need not be forced to remain monotone using constraints: if $N(x)$ and $P(x)$ are sufficiently close and $N(x)$ is strictly monotone with not too small $\min_x \frac{\partial N(x)}{\partial x}$, then we will get monotonicity of $P(x)$. Next, we point out that this formulation has a problem in approximation of any target function by a pseudo network. We decompose $P(x)$ into two parts: $P(x) = P_c(x) + P_\ell(x)$, where

$$P_c(x) = \sum_{r=1}^m a_{r0} (\tanh(w_{r0}x + b_{r0})) \quad \text{and} \quad P_\ell(x) = \sum_{r=1}^m a_{r0} (\tanh'(w_{r0}x + b_{r0})(w_r x + b_r)).$$

Note that $P_c(x)$ only depends upon initialization and does not depend on w_r and b_r . Hence, it can not approximate the target function after the training therefore, $P_\ell(x)$ needs to approximate target function with $P_c(x)$ subtracted. Now, we will show that $P_\ell(x)$ can not approximate “sufficiently non-linear” function. Note that by the initialization, $|w_{r0}|$ and $|b_{r0}|$ are $O\left(\frac{\sqrt{\log m}}{\sqrt{m}}\right)$ with high probability; therefore, $|w_{r0}x + b_{r0}|$ is $O\left(\frac{\sqrt{\log m}}{\sqrt{m}}\right)$. Using the fact that $\tanh'(y) \approx 1$ for small y , we get that $\tanh'(w_{r0}x + b_{r0}) \approx 1$ for sufficient large m . In such cases, $P_\ell(x)$ becomes linear function in x and won’t be able to approximate sufficiently non-linear function.

Note that this issue does not arise in pseudo network with ReLU activation because of discontinuity of derivative of ReLU activation at 0 but as described earlier, constrained normalizing flow need to have continuous derivative. The same issue in approximation will arise for all continuous activation function with constant second derivative. A generalization argument for activations with continuous derivatives is not known even in the supervised case, therefore we do not work with constrained normalizing flow. However, we show the effect of overparameterization for constrained normalizing flow with tanh activation in experiments (Section 3).

2.2 UNCONSTRAINED NORMALIZING FLOW

Unlike the constrained case, where we modeled $f(x)$ using a neural network $N(x)$, here we model $f'(x)$ using a neural network. Then we have $f(x) = \int_{-1}^x f'(u) du$. While this cannot be computed exactly, good approximation can be obtained via numerical integration also known as numerical quadrature of $f'(x)$. The strict monotonicity of f is achieved by ensuring that $f'(x)$ is always positive. To this end a suitable nonlinearity is applied on top of the neural network: $f'(x) = \phi(N(x))$, where $N(x)$ is as in (4) with $\rho = \sigma = \text{ReLU}$, and ϕ is the function $\text{ELU} + 1$ given by $\phi(x) = e^x \mathbb{I}[x < 0] + (x + 1) \mathbb{I}[x \geq 0]$. Thus $\phi(x) > 0$, for all $x \in \mathbb{R}$, which means that $f'(x) > 0$ for all x . Although this was the only property of $\text{ELU} + 1$ mentioned by Wehenkel & Louppe (2019), it turns out to have several other properties which we will exploit in our proof: it is 1-Lipschitz monotone increasing; its derivative is bounded from above by 1.

We denote by $\tilde{f}(x)$ the estimate of $f(x) = \int_{-1}^x f'(u) du$ obtained from $f'(x)$ via quadrature $\tilde{f}(x) = \sum_{i=1}^Q q_i f'(\tau_i(x))$. Here Q is the number of quadrature points $\tau_1(x), \dots, \tau_Q(x)$, and the $q_1, \dots, q_Q \in \mathbb{R}$ are the corresponding coefficients. Wehenkel & Louppe (2019) use Clenshaw–Curtis quadrature where the coefficients q_i can be negative.

We will use simple rectangle quadrature, which arises in Riemann integration, and uses only positive coefficients: $\tilde{f}(x) = \Delta_x [f'(-1 + \Delta_x) + f'(-1 + 2\Delta_x) \dots + f'(x)]$, where $\Delta_x = \frac{x+1}{Q}$. It is known (see e.g. Chapter 5 in Atkinson (1989) for related results) that

$$\left| \tilde{f}(x) - f(x) \right| \leq \frac{M''(x+1)^2}{2Q}, \quad \text{where } M'' = \max_{u \in [-1, x]} |f''(u)|.$$

Compared to Clenshaw–Curtis quadrature, the rectangle quadrature requires more points for similar accuracy (in our experiments this was about double). However, we use it because all the coefficients are positive, which as we will see, helps make the problem of minimizing the loss a convex optimization problem.

Instead of using f , to which we do not have access, we use \tilde{f} in the loss function, denoting it $\hat{L}(f', x)$ for the standard exponential as the base distribution to write $\hat{L}(f', x) = \tilde{f}(x) - \log f'(x)$ and $\hat{L}(f', S) = \frac{1}{n} \sum_{x \in S} \hat{L}(f', x)$. The loss $\hat{L}_G(f', x)$ for the standard Gaussian as the base distribution is defined similarly.

Let X be a random variable with density supported on $[-1, 1]$. Let the base distribution be the standard exponential, and so Z will be a random variable with the standard exponential distribution. And let $F^* : \mathbb{R} \rightarrow \mathbb{R}$ be continuous monotone increasing such that $F^{*-1}(Z)$ has the same distribution as X . Let $S = \{x_1, \dots, x_n\}$ be a set of i.i.d. samples of X . Following Allen-Zhu et al. (2019), we initialize $a_{r0} \sim \mathcal{N}(0, \epsilon_a^2)$, $w_{r0} \sim \mathcal{N}(0, \frac{1}{m})$ and $b_{r0} \sim \mathcal{N}(0, \frac{1}{m})$, where $\epsilon_a > 0$ is a small constant to be set later. The SGD updates are given by $\theta^{t+1} = \theta^t - \eta \nabla_{\theta} \hat{L}(f'_t, x^t)$ where $f'_t(x) = \phi(N_t(x))$, and $x^t \in S$ is chosen uniformly at random at each step. We can now state our main result.

Theorem 1 (informal statement of Theorem 2). *(loss function is close to optimal) For any $\epsilon > 0$, hidden layer size $m \geq \frac{C_1(F^{*'})}{\epsilon^2}$, the number of samples $n \geq \frac{C_2(F^{*'})}{\epsilon^2}$ and the number of quadrature points $Q \geq \frac{C_3(F^{*'})}{\epsilon}$, where $C_1(\cdot), C_2(\cdot), C_3(\cdot)$ are complexity measures, with probability at least 0.9, we have*

$$\mathbb{E}_{\text{sgd}} \left[\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_{x \sim \mathcal{D}} L(f'_t, x) \right] - \mathbb{E}_{x \sim \mathcal{D}} [L(F^{*'}, x)] = O(\epsilon).$$

Now recall that $\text{KL}(p_{F^*, Z} \| p_{f_t, Z}) = \mathbb{E}_X \log \frac{p_{F^*, Z}(X)}{p_{f_t, Z}(X)}$, which gives $\mathbb{E}_{\text{sgd}} \left[\frac{1}{T} \sum_{t=0}^{T-1} \text{KL}(p_{F^*, Z} \| p_{f_t, Z}) \right] = O(\epsilon)$. Recall that $p_{f, Z}(x)$ is the probability density of $f^{-1}(Z)$. Using Pinsker’s inequality, we can also bound the total variation distance between the learned and data distributions $p_{f_t, Z}$ and $p_{F^*, Z}$.

Define pseudo network $g'(x)$, which acts as proxy for $f'(x)$, as $g'(x) = \phi(P(x))$. Note that our definition of pseudo network is not the most straightforward version: $g'(x)$ is not a linear approximation of $f'(x)$. As in Allen-Zhu et al. (2019), we begin by showing the existence of a pseudo network close to the target function. However, for this we cannot use the approximation lemma in Allen-Zhu et al. (2019) as it seems to require dimension at least 2. We use the recent result of Ji et al. (2020) instead (Lemmas 1). The presence of both f' and \tilde{f} and other differences in the loss function leads to new difficulties in the analysis compared to the supervised case. We refer to the full proof due to the lack of space.

3 EXPERIMENTS

Full details of experimental setup and additional results on constrained normalizing flow as well as results on unconstrained normalizing flow is given in appendix E.

3.1 RESULTS FOR CONSTRAINED NORMALIZING FLOW

In Sec. 2.1, we suggested that high overparameterization may adversely affect training for constrained normalizing flows. We now give experimental evidence for this in Figs. 1. In Figs. 1, we see that as we increase the learning rate, training becomes more stable for larger m . Note that for learning

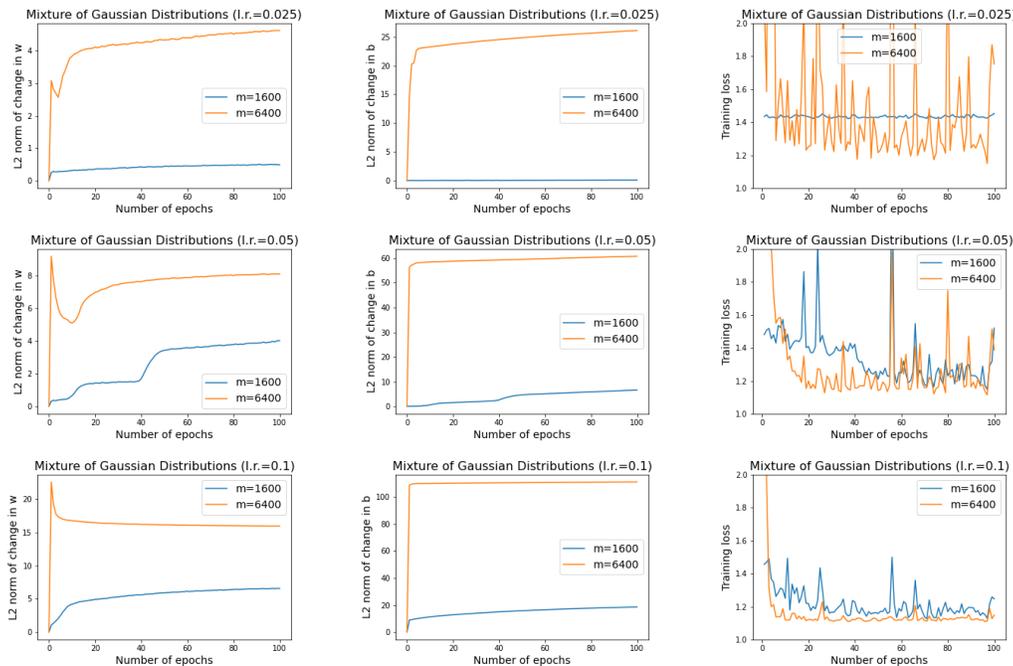


Figure 1: Effect of over-parameterization on training of constrained normalizing flow on mixture of Gaussian dataset for number of hidden layers $m = 1600, 6400$

rate equal to 0.025, constrained normalizing flow with $m = 1600$ doesn't learn anything due to small learning rate. We observe that the $L2$ -norms of W^t and B^t for $m = 6400$ are at least as large as those of $m = 1600$. On both datasets, as we increase the learning rate, $L2$ -norm of B^t increases and learning of constrained normalizing flow becomes more and more stable. These observations support our claim in Sec.2.1 that for learning and approximation of overparameterized constrained normalizing flow, neural networks need large $L2$ -norms of W^t and B^t .

4 CONCLUSION

In this paper, we gave the first theoretical analysis of normalizing flows in the simple but instructive univariate case. We gave empirical and theoretical evidence that overparametrized networks are unlikely to be useful for CNFs. By contrast, for UNFs, overparametrization does not hurt and we can adapt techniques from supervised learning to analyze two-layer (or one hidden layer) networks. Our technical adaptations and NF variants may find use in future work.

Our work raises a number of open problems: (1) We made two changes to the unconstrained flow architecture of Wehenkel & Louppe (2019). An obvious open problem is an analysis of the original architecture or with at most one change. While the exponential distribution works well as the base distribution, can we also analyze the Gaussian distribution? Similarly, Clenshaw-Curtis quadrature instead of simple rectangle quadrature? These problems seem tractable but also likely to require interesting new techniques as the optimization becomes non-convex. That would get us one step closer to the architectures used in practice. (2) Analysis of constrained normalizing flows. It is likely to be difficult because, as our results suggest, one needs networks that are not highly overparametrized—this regime is not well-understood even in the supervised case. (3) Finally, analysis of normalizing flows for the multidimensional case. Our 1D result brings into focus potential difficulties: All unconstrained architectures seem to require more than one hidden layer, which poses difficult challenges even in the supervised case. For CNFs, it is possible to design an architecture with one hidden layer, but as we have seen in our analysis of UNFs, that is challenging too.

REFERENCES

- Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. In *Advances in neural information processing systems*, pp. 6158–6169, 2019.
- Sanjeev Arora, Andrej Risteski, and Yi Zhang. Do GANs learn the distribution? some theory and empirics. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=BJehNfW0->.
- Sanjeev Arora, Simon S. Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pp. 322–332. PMLR, 2019. URL <http://proceedings.mlr.press/v97/arora19a.html>.
- Kendall Atkinson. *An Introduction to Numerical Analysis*. John Wiley and Sons, 1989.
- Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. Generating sentences from a continuous space. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pp. 10–21, Berlin, Germany, August 2016. Association for Computational Linguistics. doi: 10.18653/v1/K16-1002. URL <https://www.aclweb.org/anthology/K16-1002>.
- Rares-Darius Buhai, Andrej Risteski, Yoni Halpern, and David Sontag. Empirical study of benefits of overparameterization in single-layer latent variable generative models. In *Proceedings of the 37th International Conference on Machine Learning*, 2020. URL https://proceedings.icml.cc/static/paper_files/icml/2020/5645-Paper.pdf.
- Nicola De Cao, Wilker Aziz, and Ivan Titov. Block neural autoregressive flow. In *Proceedings of the Thirty-Fifth Conference on Uncertainty in Artificial Intelligence, UAI 2019, Tel Aviv, Israel, July 22-25, 2019*, pp. 511. AUAI Press, 2019. URL <http://auai.org/uai2019/proceedings/papers/511.pdf>.
- I. Daubechies, R. DeVore, S. Foucart, B. Hanin, and G. Petrova. Nonlinear approximation and (deep) relu networks. *ArXiv*, abs/1905.02199, 2019.
- Simon S Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. In *Proceedings of the 35th International Conference on Learning Representations*, 2018. URL <https://arxiv.org/abs/1810.02054>.
- Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In Zoubin Ghahramani, Max Welling, Corinna Cortes, Neil D. Lawrence, and Kilian Q. Weinberger (eds.), *Advances in Neural Information Processing Systems 27: Annual Conference on Neural Information Processing Systems 2014, December 8-13 2014, Montreal, Quebec, Canada*, pp. 2672–2680, 2014. URL <http://papers.nips.cc/paper/5423-generative-adversarial-nets>.
- Chin-Wei Huang, David Krueger, Alexandre Lacoste, and Aaron C. Courville. Neural autoregressive flows. In Jennifer G. Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2083–2092. PMLR, 2018. URL <http://proceedings.mlr.press/v80/huang18d.html>.
- Arthur Jacot, Clément Hongler, and Franck Gabriel. Neural tangent kernel: Convergence and generalization in neural networks. In Samy Bengio, Hanna M. Wallach, Hugo Larochelle, Kristen Grauman, Nicolò Cesa-Bianchi, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, 3-8 December 2018, Montréal, Canada*, pp. 8580–8589, 2018. URL <http://papers.nips.cc/paper/8076-neural-tangent-kernel-convergence-and-generalization-in-neural-networks>.

- Ziwei Ji, Matus Telgarsky, and Ruicheng Xian. Neural tangent kernels, transportation mappings, and universal approximation. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*, 2020. URL <https://arxiv.org/abs/1910.06956>.
- Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. In Yoshua Bengio and Yann LeCun (eds.), *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*, 2014. URL <http://arxiv.org/abs/1312.6114>.
- I. Kobyzev, S. Prince, and M. Brubaker. Normalizing flows: An introduction and review of current methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- Zhifeng Kong and Kamalika Chaudhuri. The expressive power of a class of normalizing flow models. volume 108 of *Proceedings of Machine Learning Research*, pp. 3599–3609, Online, 26–28 Aug 2020. PMLR. URL <http://proceedings.mlr.press/v108/kong20a.html>.
- Qi Lei, Jason D. Lee, Alexandros G. Dimakis, and Constantinos Daskalakis. SGD learns one-layer networks in WGANs. In *In Proceedings of the 37th International Conference on Machine Learning*, 2020. URL https://proceedings.icml.cc/static/paper_files/icml/2020/4998-Paper.pdf.
- Yuan-Chuan Li and Cheh-Chih Yeh. Some equivalent forms of bernoulli’s inequality: A survey. *Applied Mathematics*, 4(07):1070, 2013.
- Yuanzhi Li and Zehao Dou. Making method of moments great again? – how can GANs learn the target distribution, 2020. URL <https://arxiv.org/abs/2003.04033>.
- Yuanzhi Li and Yingyu Liang. Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems 31*, pp. 8157–8166, 2018. URL <http://papers.nips.cc/paper/8038-learning-overparameterized-neural-networks-via-stochastic-gradient-descent-on-structured-data.pdf>.
- Mario Lucic, Karol Kurach, Marcin Michalski, Olivier Bousquet, and Sylvain Gelly. Are gans created equal? a large-scale study. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, pp. 698–707, 2018.
- Andreas Maurer. A vector-contraction inequality for rademacher complexities. In *International Conference on Algorithmic Learning Theory*, pp. 3–17. Springer, 2016.
- Thanh V. Nguyen, Raymond K. W. Wong, and Chinmay Hegde. Benefits of jointly training autoencoders: An improved neural tangent kernel analysis. *CoRR*, abs/1911.11983, 2019a. URL <http://arxiv.org/abs/1911.11983>.
- Thanh V. Nguyen, Raymond K. W. Wong, and Chinmay Hegde. On the dynamics of gradient descent for autoencoders. In *The 22nd International Conference on Artificial Intelligence and Statistics, AISTATS 2019, 16-18 April 2019, Naha, Okinawa, Japan*, volume 89 of *Proceedings of Machine Learning Research*, pp. 2858–2867. PMLR, 2019b. URL <http://proceedings.mlr.press/v89/nguyen19a.html>.
- George Papamakarios, Theo Pavlakou, and Iain Murray. Masked autoregressive flow for density estimation. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, pp. 2335–2344, 2017. ISBN 9781510860964.
- George Papamakarios, Eric T. Nalisnick, Danilo Jimenez Rezende, Shakir Mohamed, and Balaji Lakshminarayanan. Normalizing flows for probabilistic modeling and inference. *ArXiv*, abs/1912.02762, 2019.
- Adityanarayanan Radhakrishnan, Mikhail Belkin, and Caroline Uhler. Overparameterized neural networks can implement associative memory, 2020. URL <https://arxiv.org/abs/1909.12362>.

- Danilo Rezende and Shakir Mohamed. Variational inference with normalizing flows. volume 37 of *Proceedings of Machine Learning Research*, pp. 1530–1538, Lille, France, 07–09 Jul 2015. PMLR. URL <http://proceedings.mlr.press/v37/rezende15.html>.
- Justin Romberg. Maximum of a sequence of gaussian random variables. 2012. URL <http://cnx.org/contents/8bd316d8-6442-4f5a-a597-aef1d6202f87@1>.
- Justin Sahs, Ryan Pyle, Aneel Damaraju, Josue Ortega Caro, Onur Tavaslioglu, Andy Lu, and Ankit Patel. Shallow univariate relu networks as splines: Initialization, loss surface, hessian, gradient flow dynamics, 2020.
- Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, Xi Chen, and Xi Chen. Improved techniques for training gans. In *Advances in Neural Information Processing Systems*, pp. 2234–2242. 2016. URL <http://papers.nips.cc/paper/6125-improved-techniques-for-training-gans.pdf>.
- Filippo Santambrogio. *Optimal Transport for Applied Mathematicians. Calculus of Variations, PDEs and Modeling*. Birkhäuser, 2015.
- Pedro Savarese, Itay Evron, Daniel Soudry, and Nathan Srebro. How do infinite width bounded norm networks look in function space? In Alina Beygelzimer and Daniel Hsu (eds.), *Conference on Learning Theory, COLT 2019, 25-28 June 2019, Phoenix, AZ, USA*, volume 99 of *Proceedings of Machine Learning Research*, pp. 2667–2690. PMLR, 2019. URL <http://proceedings.mlr.press/v99/savarese19a.html>.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge university press, 2014.
- Takeshi Teshima, I. Ishikawa, Koichi Tojo, Kenta Oono, M. Ikeda, and M. Sugiyama. Coupling-based invertible neural networks are universal diffeomorphism approximators. *ArXiv*, abs/2006.11469, 2020.
- Jakub M Tomczak and Max Welling. Improving variational auto-encoders using householder flow. *arXiv preprint arXiv:1611.09630*, 2016.
- Rianne van den Berg, Leonard Hasenclever, Jakub Tomczak, and Max Welling. Sylvester normalizing flows for variational inference. In *proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI)*, 2018.
- Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019. doi: 10.1017/9781108627771.
- Antoine Wehenkel and Gilles Louppe. Unconstrained monotonic neural networks. In Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d’Alché-Buc, Emily B. Fox, and Roman Garnett (eds.), *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, 8-14 December 2019, Vancouver, BC, Canada*, pp. 1543–1553, 2019. URL <http://papers.nips.cc/paper/8433-unconstrained-monotonic-neural-networks>.
- Francis Williams, Matthew Trager, Daniele Panozzo, Claudio Silva, Denis Zorin, and Joan Bruna. Gradient dynamics of shallow univariate relu networks. In *Advances in Neural Information Processing Systems 32*, pp. 8378–8387. 2019. URL <http://papers.nips.cc/paper/9046-gradient-dynamics-of-shallow-univariate-relu-networks.pdf>.
- Gilad Yehudai and Ohad Shamir. On the power and limitations of random features for understanding neural networks. In *Advances in Neural Information Processing Systems*, pp. 6598–6608, 2019.

A EXISTENCE

This section contains a proof that shows existence of a pseudo network whose loss closely approximates the loss of the target function.

Lemma 1. *For every positive function $F^{*'}$, for every x in the radius of 1 (i.e. $|x| \leq 1$), there exist a function $h(w_{r0}, b_{r0}) : \mathbb{R}^2 \rightarrow [-U_h, U_h]$ such that*

$$|\phi^{-1}(F^{*'}(x)) - \mathbb{E}_{w_{r0}, b_{r0} \sim \mathcal{N}(0,1)} [h(w_{r0}, b_{r0}) \mathbb{I}[w_{r0}x + b_{r0}]]| \leq \omega_{\phi^{-1}(F^{*'})}(\delta)$$

where U_h is given by

$$U_h = \tilde{O} \left(\frac{\|(\phi^{-1}(F^{*'}))_{|\delta}\|_{L_1}^5}{\delta^{10}(\omega_{\phi^{-1}(F^{*'})}(\delta))^4} \right)$$

Proof. We use a result from Ji et al. (2020) to prove the lemma.

Result 1. (Theorem 4.3 from Ji et al. (2020)) *Let $\psi : \mathbb{R} \rightarrow \mathbb{R}$ and $\delta > 0$ be given, and define*

$$\omega_\psi(\delta) = \sup\{\psi(x) - \psi(x') : \max\{|x|, |x'|\} \leq 1 + \delta, |x - x'| \leq \delta\}$$

$$\psi_{|\delta}(x) := \psi(x) \mathbb{I}[|x| \leq 1 + \delta]$$

$$\psi_{|\delta, \alpha} := \psi_{|\delta} * G_\alpha$$

$$\alpha := \frac{\delta}{1 + \sqrt{2 \log(2M/\omega_\psi(\delta))}} = \tilde{O}(\delta)$$

$$\beta := \frac{1}{2\pi\alpha^2}$$

$$\begin{aligned} \mathcal{T}_r(w_{r0}, b_{r0}) := & 2 \left[\psi_{|\delta, \alpha}(0) + \int \left| \hat{\psi}_{|\delta, \alpha}(v) \right| \cos(2\pi(\theta_{\psi_{|\delta, \alpha}}(v) - \|v\|)) dv \right] \\ & + 2\pi(2\pi\beta^2) \left| \hat{\psi}_{|\delta}(\beta w_{r0}) \right| e^{\frac{(b_{r0})^2}{2}} \sin(2\pi(\theta_{\psi_{|\delta, \alpha}}(\beta w_{r0}) - b_{r0})) \mathbb{I}[|b_{r0}| \leq \|w_{r0}\| \leq r] \end{aligned}$$

where $*$ denotes convolution operation, G_α denotes Gaussian with coordinate-wise variance α^2 .

$\left| \hat{\psi}_{|\delta, \alpha} \right|$ denotes magnitude of fourier transform of $\psi_{|\delta, \alpha}$ and $\theta_{\psi_{|\delta, \alpha}}$ denotes phase of fourier transform.

Then,

$$\sup_{|x| \leq 1} |\psi(x) - \mathbb{E}_{w_{r0}, b_{r0} \sim \mathcal{N}(0,1)} [\mathcal{T}_r(w_{r0}, b_{r0}) \mathbb{I}[w_{r0}x + b_{r0} \geq 0]]| \leq \omega_\psi(\delta) \quad (6)$$

The upper bound of $\mathcal{T}_r(w_{r0}, b_{r0})$ is given by

$$\sup_{w_{r0}, b_{r0}} \|\mathcal{T}_r(w_{r0}, b_{r0})\| = \tilde{O} \left(\frac{\|\psi_{|\delta}\|_{L_1}^5}{\delta^{10}(\omega_\psi(\delta))^4} \right) = U_{\mathcal{T}} \quad (7)$$

Using result 1 for $\phi^{-1}(F^{*'}(x))$ function, denoting $\mathcal{T}_r(w_{r0}, b_{r0})$ for $\phi^{-1}(F^{*'}(x))$ function as $h(w_{r0}, b_{r0})$, we get

$$|\phi^{-1}(F^{*'}(x)) - \mathbb{E}_{w_{r0}, b_{r0} \sim \mathcal{N}(0,1)} [h(w_{r0}, b_{r0}) \mathbb{I}[w_{r0}x + b_{r0}]]| \leq \omega_{\phi^{-1}(F^{*'})}(\delta)$$

with following upper bound on $h(w_{r0}, b_{r0})$.

$$\sup_{w_{r0}, b_{r0}} h(w_{r0}, b_{r0}) \leq \tilde{O} \left(\frac{\|(\phi^{-1}(F^{*'}))_{|\delta}\|_{L_1}^5}{\delta^{10}(\omega_{\phi^{-1}(F^{*'})}(\delta))^4} \right) = U_h$$

□

Write pseudo network $P(x)$ as

$$P(x) = P_c(x) + P_\ell(x)$$

where

$$P_c(x) = \sum_{r=1}^m a_{r0} (w_{r0}x + b_{r0}) \mathbb{I}[w_{r0}x + b_{r0}]$$

$$P_\ell(x) = \sum_{r=1}^m a_{r0} (w_r x + b_r) \mathbb{I}[w_{r0}x + b_{r0}]$$

Lemma 2. (Approximating target function using linear pseudo network) For every positive function $F^{*'}$, with at least $1 - \frac{1}{c_1} - \exp\left(-\frac{\epsilon^2 m}{128c_1^2 U_h^2 \log m}\right)$ probability over random initialization, for all $x \in [-1, 1]$ we get

$$|\phi(P_\ell^*(x)) - F^{*'}(x)| \leq \omega_{\phi^{-1}(F^{*'})}(\delta) + \epsilon$$

and upper bound L_∞ norm of parameters is given by

$$\|\theta^*\|_\infty \leq \frac{U_h \sqrt{\pi}}{\sqrt{2m\epsilon_a}}$$

Proof. Define w_r^* and b_r^* as

$$w_r^* = 0$$

$$b_r^* = \frac{\text{sign}(a_{r0}) \sqrt{\pi}}{m\epsilon_a \sqrt{2}} h(\sqrt{m}w_{r0}, \sqrt{m}b_{r0}) \quad (8)$$

Using w_r^* and b_r^* ,

$$\begin{aligned} & \mathbb{E}_{a_{r0} \sim \mathcal{N}(0, \epsilon_a^2), w_{r0} \sim \mathcal{N}(0, \frac{1}{m}), b_{r0} \sim \mathcal{N}(0, \frac{1}{m})} [P_\ell^*(x)] \\ &= \mathbb{E}_{a_{r0} \sim \mathcal{N}(0, \epsilon_a^2), w_{r0} \sim \mathcal{N}(0, \frac{1}{m}), b_{r0} \sim \mathcal{N}(0, \frac{1}{m})} \left[\sum_{r=1}^m a_{r0} (w_r^* x + b_r^*) \mathbb{I}[w_{r0}x + b_{r0} \geq 0] \right] \\ &= \mathbb{E}_{a_{r0} \sim \mathcal{N}(0, \epsilon_a^2), w_{r0} \sim \mathcal{N}(0, \frac{1}{m}), b_{r0} \sim \mathcal{N}(0, \frac{1}{m})} \left[\frac{a_{r0} \text{sign}(a_{r0}) \sqrt{\pi}}{\epsilon_a \sqrt{2}} h(\sqrt{m}w_{r0}, \sqrt{m}b_{r0}) \mathbb{I}[w_{r0}x + b_{r0} \geq 0] \right] \\ &\stackrel{(i)}{=} \mathbb{E}_{w_{r0} \sim \mathcal{N}(0, \frac{1}{m}), b_{r0} \sim \mathcal{N}(0, \frac{1}{m})} [h(\sqrt{m}w_{r0}, \sqrt{m}b_{r0}) \mathbb{I}[\sqrt{m}(w_{r0}x + b_{r0}) \geq 0]] \end{aligned}$$

where equality (i) follows from $\mathbb{E}_{a_{r0} \sim \mathcal{N}(0, \epsilon_a^2)} [a_{r0} \text{sign}(a_{r0})] = \epsilon_a \sqrt{\frac{2}{\pi}}$ and homogeneity of indicator function. Using lemma 1,

$$\begin{aligned} & \left| \mathbb{E}_{a_{r0} \sim \mathcal{N}(0, \epsilon_a^2), w_{r0} \sim \mathcal{N}(0, \frac{1}{m}), b_{r0} \sim \mathcal{N}(0, \frac{1}{m})} [P_\ell^*(x)] - \phi^{-1}(F^{*'}(x)) \right| \\ &= \left| \mathbb{E}_{w_{r0} \sim \mathcal{N}(0, \frac{1}{m}), b_{r0} \sim \mathcal{N}(0, \frac{1}{m})} [h(\sqrt{m}w_{r0}, \sqrt{m}b_{r0}) \mathbb{I}[\sqrt{m}(w_{r0}x + b_{r0}) \geq 0]] - \phi^{-1}(F^{*'}(x)) \right| \\ &\leq \omega_{\phi^{-1}(F^{*'})}(\delta) \quad (9) \end{aligned}$$

Using technique from Yehudai & Shamir (2019), we define

$$h = h((a_{10}, w_{10}, b_{10}), \dots, (a_{r0}, w_{r0}, b_{r0}), \dots, (a_{m0}, w_{m0}, b_{m0})) = \sup_{x \in [-1, 1]} |P_\ell^*(x) - \mathbb{E}_{a_{r0}, w_{r0}, b_{r0}} [P_\ell^*(x)]|$$

We will use McDiarmid's inequality to bound h .

$$\begin{aligned} & \left| h((a_{10}, w_{10}, b_{10}), \dots, (a_{r0}, w_{r0}, b_{r0}), \dots, (a_{10}, w_{m0}, b_{m0})) - h((a'_{10}, w'_{10}, b'_{10}), \dots, (a'_{r0}, w'_{r0}, b'_{r0}), \dots, (a_{10}, w_{m0}, b_{m0})) \right| \\ &\leq \frac{4c_1 U_h \sqrt{2 \log m}}{m} \end{aligned}$$

Using Lemma 26.2 from Shalev-Shwartz & Ben-David (2014), we get

$$\mathbb{E}[h] = \frac{2}{m} \mathbb{E}_{a_{r0}, w_{r0}, b_{r0}, \xi_r} \left[\sup_x m \left| \sum_{r=1}^m \xi_i (w_r^* x + b_r^*) \mathbb{I}[w_{r0}x + b_{r0}] \right| \right]$$

where $\xi_1, \xi_2, \dots, \xi_m$ are independent Rademacher random variables.

$$\begin{aligned} \mathbb{E}_{a_{r0}, w_{r0}, b_{r0}} [h] &\leq \frac{2}{m} \mathbb{E}_{a_{r0}, w_{r0}, b_{r0}, \xi_r} \left[\sup_x m \left| \sum_{r=1}^m \xi_i a_{r0} (w_r^* x + b_r^*) \mathbb{I}[w_{r0} x + b_{r0}] \right| \right] \\ &\leq \frac{2}{m} \mathbb{E}_{a_{r0}, w_{r0}, b_{r0}, \xi_r} \left[\sup_x m \left| \sum_{r=1}^m \xi_i a_{r0} (w_r^* x + b_r^*) \mathbb{I}[w_{r0} x + b_{r0}] \right| \right] \\ &\leq \frac{8c_1 \sqrt{\log m} U_h}{m} \mathbb{E}_{a_{r0}, w_{r0}, b_{r0}, \xi_r} \left[\sup_x \left| \sum_{r=1}^m \xi_i \mathbb{I}[w_{r0} x + b_{r0}] \right| \right] \end{aligned}$$

One can show that

$$\frac{1}{m} \mathbb{E}_{a_{r0}, w_{r0}, b_{r0}, \xi_r} \left[\sup_x \left| \sum_{r=1}^m \xi_i \mathbb{I}[w_{r0} x + b_{r0}] \right| \right] \leq 2\sqrt{\frac{\log m}{m}}$$

Using this relation, we get

$$\mathbb{E}_{a_{r0}, w_{r0}, b_{r0}} [h] \leq \frac{16c_1 U_h \log m}{\sqrt{m}}$$

Using Mcdiarmid's inequality, with at least $1 - \frac{1}{c_1} - \exp\left(-\frac{\epsilon^2 m}{128c_1^2 U_h^2 \log m}\right)$, we have

$$|P_\ell^*(x) - \mathbb{E}_{a_{r0}, w_{r0}, b_{r0}} [P_\ell^*(x)]| = h \leq \frac{\epsilon}{2} + \frac{16c_1 U_h \log m}{\sqrt{m}} \stackrel{(i)}{\leq} \epsilon \quad (10)$$

where inequality (i) follows from our choice of m in lemma 10. Using eq.(9), we get

$$|P_\ell^*(x) - \phi^{-1}(F^{*'}(x))| \leq \omega_{\phi^{-1}(F^{*'})}(\delta) + \epsilon$$

Using 1-Lipschitzness of ϕ , we get

$$\begin{aligned} |\phi(P_\ell^*(x)) - F^{*'}(x)| &= |\phi(P_\ell^*(x)) - \phi(\phi^{-1}(F^{*'}(x)))| \\ &\leq |P_\ell^*(x) - \phi^{-1}(F^{*'}(x))| \\ &\leq \omega_{\phi^{-1}(F^{*'})}(\delta) + \epsilon \end{aligned}$$

The upper bound on norm of $\|\theta^*\|_\infty$ is given by the following equation.

$$\|\theta^*\|_\infty \leq \frac{U_h \sqrt{\pi}}{\sqrt{2m\epsilon_a}}$$

□

Corollary 1. (Approximating target network using complete pseudo network) For every positive function $F^{*'}$, with at least $0.99 - \frac{1}{c_1} - \frac{1}{c_6} - \frac{1}{c_7} - \exp\left(-\frac{\epsilon^2 m}{128c_1^2 U_h^2 \log m}\right)$ probability

$$|\phi(P^*(x)) - F^{*'}(x)| \leq 16c_1 (c_6 + c_7) \epsilon_a \log m + \omega_{\phi^{-1}(F^{*'})}(\delta) + \epsilon$$

and upper bound L_∞ norm of parameters is given by

$$\|\theta^*\|_\infty \leq \frac{U_h \sqrt{\pi}}{\sqrt{2m\epsilon_a}}$$

Proof. Using Lipschitz continuity of ϕ function, we get

$$\begin{aligned} |\phi(P_\ell^*(x)) - \phi(P^*(x))| &\leq |P_\ell^*(x) - P^*(x)| \\ &\leq \left| \sum_{r=1}^m a_{r0} (w_{r0} x + b_{r0}) \mathbb{I}[w_{r0} x + b_{r0} \geq 0] \right| \end{aligned}$$

Now, there are at most m break points of indicator $\mathbb{I}[w_{r0} x + b_{r0} \geq 0]$ where value of $\mathbb{I}[w_{r0} x + b_{r0} \geq 0]$ changes. We can divide range of x into at most $m + 1$ subsets where in each

subset, value of indicators $\mathbb{I}[w_{r0}x + b_{r0} \geq 0]$ is fixed for all r . Suppose there are m' indicators with value 1 in a given subset. Without loss of generality, we can assume that indicators from $r = 1$ to $r = m'$ is 1. Then,

$$\begin{aligned} \left| \sum_{r=1}^m a_{r0} (w_{r0}x + b_{r0}) \mathbb{I}[w_{r0}x + b_{r0} \geq 0] \right| &= \left| \sum_{r=1}^{m'} a_{r0} (w_{r0}x + b_{r0}) \right| \\ &\leq \left| x \sum_{r=1}^{m'} a_{r0} w_{r0} + \sum_{r=1}^{m'} a_{r0} b_{r0} \right| \end{aligned}$$

Now, applying Hoeffding's inequality for the sum in above equation, we get

$$\begin{aligned} \Pr \left(\left| \sum_{r=1}^{m'} a_{r0} w_{r0} \right| \geq t \right) &\leq \exp \left(- \frac{2t^2 m}{m' (2c_1 \epsilon_a \sqrt{2 \log m})^2 (2c_6 \sqrt{2 \log m})^2} \right) \\ &= \exp \left(- \frac{t^2}{32c_1^2 c_6^2 \epsilon_a^2 (\log m)^2} \right) \end{aligned}$$

Taking $t = 16c_1 c_6 \epsilon_a (\log m)$, with at least probability $0.999 - \frac{1}{c_1} - \frac{1}{c_6}$, we have

$$\left| \sum_{r=1}^{m'} a_{r0} w_{r0} \right| \leq 16c_1 c_6 \epsilon_a (\log m)$$

and similarly, we will get that with at least $0.999 - \frac{1}{c_1} - \frac{1}{c_7}$ probability,

$$\left| \sum_{r=1}^{m'} a_{r0} b_{r0} \right| \leq 16c_1 c_7 \epsilon_a (\log m)$$

we will get that at least $0.999 - \frac{1}{c_1} - \frac{1}{c_6} - \frac{1}{c_7}$ probability, we have

$$\begin{aligned} \left| \sum_{r=1}^m a_{r0} w_{r0} \mathbb{I}[w_{r0}x + b_{r0} \geq 0] \right| &\leq 16c_1 c_6 \epsilon_a (\log m) \\ \left| \sum_{r=1}^m a_{r0} b_{r0} \mathbb{I}[w_{r0}x + b_{r0} \geq 0] \right| &\leq 16c_1 c_7 \epsilon_a (\log m) \end{aligned} \quad (11)$$

Using these relations, we get that with at least $0.99 - \frac{1}{c_1} - \frac{1}{c_6} - \frac{1}{c_7}$ probability,

$$\left| \sum_{r=1}^m a_{r0} (w_{r0}x + b_{r0}) \mathbb{I}[w_{r0}x + b_{r0} \geq 0] \right| \leq 16c_1 (c_6 + c_7) \epsilon_a \log m \quad (12)$$

Using above inequality, we get

$$|\phi(P_\ell^*(x)) - \phi(P^*(x))| \leq |P_\ell^*(x) - P^*(x)| \leq 16c_1 (c_6 + c_7) \epsilon_a \log m$$

Using lemma 2, with at least $0.99 - \frac{1}{c_1} - \frac{1}{c_6} - \frac{1}{c_7} - \exp\left(-\frac{\epsilon^2 m}{128c_1^2 U_h^2 \log m}\right)$ probability,

$$\begin{aligned} |\phi(P^*(x)) - F^{*'}(x)| &\leq |\phi(P^*(x)) - \phi(P_\ell^*(x))| + |\phi(P_\ell^*(x)) - F^{*'}(x)| \\ &\leq 16c_1 (c_6 + c_7) \epsilon_a \log m + \omega_{\phi^{-1}(F^{*'})}(\delta) + \epsilon \end{aligned}$$

□

Lemma 3. (Optimal loss) For every positive function $F^{*'}$, with at least $0.99 - \frac{1}{c_1} - \frac{1}{c_6} - \frac{1}{c_7} - \exp\left(-\frac{\epsilon^2 m}{128c_1^2 U_h^2 \log m}\right)$ probability, there exist a pseudo network whose loss is close to that of the target function for all $x \in [-1, 1]$.

$$\left| \hat{L}(\phi(P^*), x) - \hat{L}(F^{*'}, x) \right| \leq 3(16c_1(c_6 + c_7)\epsilon_a \log m + \omega_{\phi^{-1}(F^{*'})}(\delta) + \epsilon)$$

Proof.

$$\begin{aligned} \left| \hat{L}(\phi(P^*), x) - \hat{L}(F^{*'}, x) \right| &\leq \left| \sum_{i=1}^{\tilde{n}} \tilde{w}_i \phi(P^*(\tau_i(x))) - \sum_{i=1}^{\tilde{n}} \tilde{w}_i F^{*'}(\tau_i(x)) \right| \\ &\quad + |\log(\phi(P^*(x))) - \log(F^{*'}(x))| \\ &\leq 2(16c_1(c_6 + c_7)\epsilon_a \log m \\ &\quad + \omega_{\phi^{-1}(F^{*'})}(\delta) + \epsilon + |P^*(x) - \phi^{-1}(F^{*'}(x))|) \\ &\leq 2(16c_1(c_6 + c_7)\epsilon_a \log m \\ &\quad + \omega_{\phi^{-1}(F^{*'})}(\delta) + \epsilon + |P_c^*(x)| + |P_\ell^*(x) - \phi^{-1}(F^{*'}(x))|) \\ &\leq 3(16c_1(c_6 + c_7)\epsilon_a \log m + \omega_{\phi^{-1}(F^{*'})}(\delta) + \epsilon) \end{aligned}$$

□

B COUPLING

In this section, we prove that, for random initialization, the gradients of the loss of pseudo network closely approximate the gradients of the loss of the target function. In other words, we show coupling of their gradient-based optimizations. Define \tilde{L}_1 as

$$\tilde{L}_1 = \sup_{t \in \{1, 2, \dots, T\}, w_r^t, b_r^t, |x| \leq 1} \frac{\phi'(N_t(x))}{\phi(N_t(x))} \quad (13)$$

We get following find upper bound on \tilde{L}_1 .

$$\begin{aligned} \tilde{L}_1 &= \sup_{t \in \{1, 2, \dots, T\}, w_r^t, b_r^t, |x| \leq 1} \frac{\phi'(N_t(x))}{\phi(N_t(x))} \\ &= \sup_{t \in \{1, 2, \dots, T\}, w_r^t, b_r^t, |x| \leq 1} \frac{\exp(N_t(x)) \mathbb{I}[N_t(x) < 0] + \mathbb{I}[N_t(x) \geq 0]}{\exp(N_t(x)) \mathbb{I}[N_t(x) < 0] + (N_t(x) + 1) \mathbb{I}[N_t(x) \geq 0]} \\ &= \sup_{t \in \{1, 2, \dots, T\}, w_r^t, b_r^t, |x| \leq 1} \mathbb{I}[N_t(x) < 0] + \frac{\mathbb{I}[N_t(x) \geq 0]}{N_t(x) + 1} \\ &= 1 \end{aligned} \quad (14)$$

Define $\bar{\Delta}$ as

$$\bar{\Delta} = 6c_1 \epsilon_a \sqrt{2 \log m} \quad (15)$$

for some constant c_1 .

Lemma 4. (Bound in change in patterns) For all x in 1 radius ($|x| \leq 1$), with probability at least $1 - \frac{1}{c_1} - \exp\left(-\frac{64(1-c_2)^2 \eta^2 m^2 \bar{\Delta}^2 t^2}{\pi}\right)$ over random initialization, for at most $c_2 \frac{4\sqrt{2}\eta\sqrt{m}\bar{\Delta}t}{\sqrt{\pi}}$ fraction of $r \in [m]$

$$\mathbb{I}[(w_{r_0} + w_r^t)x + b_{r_0} + b_r^t] \neq \mathbb{I}[w_{r_0}x + b_{r_0}]$$

Proof. Taking derivative of $\hat{L}(f', x)$ wrt w_r ,

$$\begin{aligned} \left| \frac{\partial \hat{L}(f', x)}{\partial w_r} \right| &= \left| \left(\sum_{i=1}^{\tilde{n}} \tilde{w}_i \phi'(N_t(\tau_i(x))) a_{r_0} \sigma'((w_{r_0} + w_r^t)\tau_i(x) + b_{r_0} + b_r^t) \tau_i(x) \right) \right| \\ &\quad + \left| \frac{1}{\phi(N_t(x))} (\phi'(N_t(x)) a_{r_0} \sigma'((w_{r_0} + w_r^t)x + b_{r_0} + b_r^t)) \right| \end{aligned}$$

Using $|x| \leq 1$, we get

$$\left| \frac{\partial \hat{L}(f'_t, x)}{\partial w_r} \right| \leq 3 |a_{r0}|$$

Using Lemma 15, with at least $1 - \frac{1}{c_1}$ probability, we get

$$\left| \frac{\partial \hat{L}(f'_t, x)}{\partial w_r} \right| \leq \bar{\Delta} \quad (16)$$

where $\bar{\Delta}$ as

$$\bar{\Delta} = 6c_1 \epsilon_a \sqrt{2 \log m}$$

Using same procedure for b_r , we get

$$\begin{aligned} \left| \frac{\partial \hat{L}(f'_t, x)}{\partial b_r} \right| &= \left| \sum_{i=1}^{\bar{n}} \tilde{w}_i \phi'(N_t(\tau_i(x))) a_{r0} \sigma'((w_{r0} + w_r^t) \tau_i(x) + b_{r0} + b_r^t) \right| \\ &\quad + \left| \frac{1}{\phi(N_t(x))} (\phi'(N_t(x)) a_{r0} \sigma'((w_{r0} + w_r^t)x + b_{r0} + b_r^t)) \right| \\ &\leq 3 |a_{r0}| \\ &= \bar{\Delta} \end{aligned} \quad (17)$$

Using eq.(16) and eq.(17), we get

$$\begin{aligned} |w_r^t| &\leq \eta \bar{\Delta} t \\ |b_r^t| &\leq \eta \bar{\Delta} t \end{aligned} \quad (18)$$

Define

$$\mathcal{H} = \{r \in [m] \mid |w_{r0}x + b_{r0}| \geq 4\eta \bar{\Delta} t\} \quad (19)$$

For all x with $|x| \leq 1$ and for all $r \in [m]$, $|w_r^t x + b_r^t| \leq 2\eta \bar{\Delta} t$. For all $r \in \mathcal{H}$, we get $\mathbb{I}[(w_{r0} + w_r^t)x + b_{r0} + b_r^t] = \mathbb{I}[w_{r0}x + b_{r0}]$. Now, we need to bound the size of \mathcal{H} . We know that $w_{r0}x + b_{r0}$ is Gaussian with $\mathbb{E}[w_{r0}x + b_{r0}] = 0$ and $\text{Var}[w_{r0}x + b_{r0}] \geq \frac{1}{m}$. Using lemma 16, we get

$$\begin{aligned} \Pr(|w_{r0}x + b_{r0}| \leq 4\eta \bar{\Delta} t) &\leq \frac{4\sqrt{2}\eta\sqrt{m}\bar{\Delta}t}{\sqrt{\pi}} \\ \Pr(|w_{r0}x + b_{r0}| \geq 4\eta \bar{\Delta} t) &\geq 1 - \frac{4\sqrt{2}\eta\sqrt{m}\bar{\Delta}t}{\sqrt{\pi}} \end{aligned}$$

Using Fact 1 for \mathcal{H}^c (where $\mathcal{H}^c = [m]/\mathcal{H}$), we get

$$\begin{aligned} \Pr\left(\mathcal{H}^c \leq c_2 m \frac{4\sqrt{2}\eta\sqrt{m}\bar{\Delta}t}{\sqrt{\pi}}\right) &\leq \exp\left(-2m \left(1 - c_2\right) \left(\frac{4\sqrt{2}\eta\sqrt{m}\bar{\Delta}t}{\sqrt{\pi}}\right)^2\right) \\ &\leq \exp\left(-\frac{64(1 - c_2)^2 \eta^2 m^2 \bar{\Delta}^2 t^2}{\pi}\right) \\ \Pr\left(\mathcal{H} \geq m \left(1 - c_2 \frac{4\sqrt{2}\eta\sqrt{m}\bar{\Delta}t}{\sqrt{\pi}}\right)\right) &\geq 1 - \exp\left(-\frac{64(1 - c_2)^2 \eta^2 m^2 \bar{\Delta}^2 t^2}{\pi}\right) \end{aligned}$$

Hence, with high probability, we have

$$\mathcal{H} \geq m \left(1 - c_2 \frac{4\sqrt{2}\eta\sqrt{m}\bar{\Delta}t}{\sqrt{\pi}}\right)$$

□

Lemma 5. (Bound on difference of f' and g') For all x in 1 radius ($|x| \leq 1$), with probability at least $1 - \frac{1}{c_1} - \exp\left(-\frac{64(1-c_2)^2\eta^2m^2\bar{\Delta}^2t^2}{\pi}\right)$ over random initialization, function with neural network and function with pseudo network are close.

$$|\phi(N_t(x)) - \phi(P_t(x))| \leq \frac{192\eta^2m^{1.5}\bar{\Delta}^2c_1c_2\epsilon_a t^2\sqrt{\log m}}{\sqrt{\pi}}$$

Proof. We know that ϕ is 1-Lipschitz continuous. Using Lipschitz continuity of ϕ , we get

$$|\phi(N_t(x)) - \phi(P_t(x))| \leq |N_t(x) - P_t(x)|$$

Bounding $|N_t(x) - P_t(x)|$,

$$\begin{aligned} |N_t(x) - P_t(x)| &\leq \left| \sum_{r \in [m]} a_{r0} ((w_{r0} + w_r^t)x + b_{r0} + b_r^t) \mathbb{I}[(w_{r0} + w_r^t)x + b_{r0} + b_r^t] \right. \\ &\quad \left. - \sum_{r \in [m]} a_{r0} ((w_{r0} + w_r^t)x + b_{r0} + b_r^t) \mathbb{I}[w_{r0}x + b_{r0}] \right| \\ &\leq \left| \sum_{r \notin \mathcal{H}} a_{r0} ((w_{r0} + w_r^t)x + b_{r0} + b_r^t) (\mathbb{I}[(w_{r0} + w_r^t)x + b_{r0} + b_r^t] - \mathbb{I}[w_{r0}x + b_{r0}]) \right| \\ &\stackrel{(i)}{\leq} \left(\frac{4\sqrt{2}c_2\eta m^{1.5}\bar{\Delta}t}{\sqrt{\pi}} \right) (2c_1\epsilon_a\sqrt{2\log m}) (4\eta\bar{\Delta}t + 2\eta\bar{\Delta}t) \quad (2) \\ &\leq \left(\frac{192\eta m^{1.5}\bar{\Delta}c_1c_2\epsilon_a t\sqrt{\log m}}{\sqrt{\pi}} \right) (\eta\bar{\Delta}t) \\ &= \frac{192\eta^2m^{1.5}\bar{\Delta}^2c_1c_2\epsilon_a t^2\sqrt{\log m}}{\sqrt{\pi}} \quad (20) \\ &\leq O(\eta^2m^{1.5}\bar{\Delta}^2\epsilon_a t^2\sqrt{\log m}) \end{aligned}$$

where inequality (i) uses lemma 4, lemma 15, eq.(18) and eq.(19). Define Δ_{np}^t as

$$\Delta_{np}^t = \frac{192\eta^2m^{1.5}\bar{\Delta}^2c_1c_2\epsilon_a t^2\sqrt{\log m}}{\sqrt{\pi}} \quad (21)$$

□

Lemma 6. (Coupling of loss functions) For all x in 1 radius ($|x| \leq 1$), with probability at least $1 - \frac{1}{c_1} - \exp\left(-\frac{64(1-c_2)^2\eta^2m^2\bar{\Delta}^2t^2}{\pi}\right)$ over random initialization, loss function of neural network and pseudo network are close.

$$|\hat{L}(f'_t, x) - \hat{L}(g'_t, x)| \leq 3\Delta_{np}^t$$

Proof.

$$\begin{aligned} |\hat{L}(f'_t, x) - \hat{L}(g'_t, x)| &\leq \left| \sum_{i=1}^{\tilde{n}} \tilde{w}_i f'_t(\tau_i(x)) - \sum_{i=1}^{\tilde{n}} \tilde{w}_i g'_t(\tau_i(x)) \right| + |\log(f'_t(x)) - \log(g'_t(x))| \\ &\stackrel{(i)}{\leq} 2 \left(\sup_{i \in [\tilde{n}]} |f'_t(\tau_i(x)) - g'_t(\tau_i(x))| \right) + |N_t(x) - P_t(x)| \\ &\stackrel{(ii)}{\leq} 3\Delta_{np}^t \end{aligned}$$

where inequality (i) follows from 1-Lipschitz continuity of $\log(\phi(N(x)))$ with respect to $N(x)$. Inequality (ii) uses eq.(20) and lemma 5. □

Lemma 7. (Coupling of gradient of functions) For all x in 1 radius ($|x| \leq 1$), with probability at least $1 - \frac{1}{c_1} - \exp\left(-\frac{64(1-c_2)^2\eta^2m^2\bar{\Delta}^2t^2}{\pi}\right)$ over random initialization, gradient of derivative of neural network function and derivative of pseudo network function with respect to parameters are close.

$$\left\|\nabla_{\theta}f'_t(x) - \nabla_{\theta}g'_t(x)\right\|_1 \leq 4c_1\epsilon_a m\Delta_{np}^t \sqrt{2\log m} + \left(\frac{64\eta m^{1.5}\bar{\Delta}c_1c_2\epsilon_a t\sqrt{\log m}}{\sqrt{\pi}}\right)$$

Proof.

$$\begin{aligned} \left\|\nabla_{\theta}f'_t(x) - \nabla_{\theta}g'_t(x)\right\|_1 &\leq \left\|\phi'(N_t(x))\nabla_{\theta}N_t(x) - \phi'(P_t(x))\nabla_{\theta}P_t(x)\right\|_1 \\ &\leq \left\|\phi'(N_t(x))\nabla_{\theta}N_t(x) - \phi'(P_t(x))\nabla_{\theta}N_t(x)\right\|_1 + \left\|\phi'(P_t(x))\nabla_{\theta}N_t(x) - \phi'(P_t(x))\nabla_{\theta}P_t(x)\right\|_1 \\ &\leq |\phi'(N_t(x)) - \phi'(P_t(x))| \left\|\nabla_{\theta}N_t(x)\right\|_1 + |\phi'(P_t(x))| \left\|\nabla_{\theta}N_t(x) - \nabla_{\theta}P_t(x)\right\|_1 \\ &\leq |N_t(x) - P_t(x)| \left\|\nabla_{\theta}N_t(x)\right\|_1 + \left\|\nabla_{\theta}N_t(x) - \nabla_{\theta}P_t(x)\right\|_1 \end{aligned}$$

where last inequality follows from 1-Lipschitzness of ϕ' function and $\phi'(x) \leq 1$ for all x such that $|x| \leq 1, t \in [T]$. To upper bound $\left\|\nabla_{\theta}N_t(x) - \nabla_{\theta}P_t(x)\right\|_1$,

$$\begin{aligned} \left\|\nabla_{\theta}N_t(x) - \nabla_{\theta}P_t(x)\right\|_1 &\leq \left\|(A_0, A_0) \odot (x, 1) \odot (\mathbb{I}[(w_{r0} + w_r^t)x + b_{r0} + b_r^t] - \mathbb{I}[w_{r0}x + b_{r0}]), \right. \\ &\quad \left. \mathbb{I}[(w_{r0} + w_r^t)x + b_{r0} + b_r^t] - \mathbb{I}[w_{r0}x + b_{r0}]\right\|_1 \\ &\leq \left(8c_1\epsilon_a\sqrt{2\log m}\right) \left(c_2\frac{4\sqrt{2}\eta m^{1.5}\bar{\Delta}t}{\sqrt{\pi}}\right) \\ &\leq \left(\frac{64\eta m^{1.5}\bar{\Delta}c_1c_2\epsilon_a t\sqrt{\log m}}{\sqrt{\pi}}\right) \end{aligned} \tag{22}$$

Using eq.(21) and eq.(22), we get

$$\begin{aligned} \left\|\nabla_{\theta}f'_t(x) - \nabla_{\theta}g'_t(x)\right\|_1 &\leq |N_t(x) - P_t(x)| \left\|(A_0, A_0) \odot (x, 1) \odot (\mathbb{I}[(w_{r0} + w_r^t)x + b_{r0} + b_r^t], \right. \\ &\quad \left. \mathbb{I}[(w_{r0} + w_r^t)x + b_{r0} + b_r^t])\right\|_1 + \left\|(A_0, A_0) \odot (x, 1) \odot (\mathbb{I}[(w_{r0} + w_r^t)x + b_{r0} + b_r^t] \right. \\ &\quad \left. - \mathbb{I}[w_{r0}x + b_{r0}], \mathbb{I}[(w_{r0} + w_r^t)x + b_{r0} + b_r^t] - \mathbb{I}[w_{r0}x + b_{r0}])\right\|_1 \\ &\leq 4c_1\epsilon_a m\Delta_{np}^t \sqrt{2\log m} + \left(8c_1\epsilon_a\sqrt{2\log m}\right) \left(c_2\frac{4\sqrt{2}\eta m^{1.5}\bar{\Delta}t}{\sqrt{\pi}}\right) \\ &\leq 4c_1\epsilon_a m\Delta_{np}^t \sqrt{2\log m} + \left(\frac{64\eta m^{1.5}\bar{\Delta}c_1c_2\epsilon_a t\sqrt{\log m}}{\sqrt{\pi}}\right) \end{aligned}$$

□

Lemma 8. (Coupling of gradient of loss) For all x in 1 radius ($|x| \leq 1$), with probability at least $1 - \frac{1}{c_1} - \exp\left(-\frac{64(1-c_2)^2\eta^2m^2\bar{\Delta}^2t^2}{\pi}\right) - \frac{1}{c_3} - 2\exp\left(-\frac{m}{8\epsilon_a^2}\right)$ over random initialization, gradient of loss function with neural network and loss function with pseudo network are close.

$$\left\|\nabla_{\theta}\hat{L}(f'_t, x) - \nabla_{\theta}\hat{L}(g'_t, x)\right\|_1 \leq \frac{192\eta m^{1.5}\bar{\Delta}c_1c_2\epsilon_a t\sqrt{\log m}}{\sqrt{\pi}} + 16c_1\epsilon_a m\Delta_{np}^t \sqrt{2\log m}$$

Proof.

$$\begin{aligned}
\|\nabla_{\theta}\hat{L}(f'_t, x) - \nabla_{\theta}\hat{L}(g'_t, x)\|_1 &\leq \left\| \sum_{i=1}^{\tilde{n}} \tilde{w}_i \nabla_{\theta} f'_t(\tau_i(x)) - \frac{\nabla_{\theta} f'_t(x)}{f'_t(x)} \right. \\
&\quad \left. - \sum_{i=1}^{\tilde{n}} \tilde{w}_i \nabla_{\theta} g'_t(\tau_i(x)) + \frac{\nabla_{\theta} g'_t(x)}{g'_t(x)} \right\|_1 \\
&\leq \left\| \sum_{i=1}^{\tilde{n}} \tilde{w}_i \nabla_{\theta} f'_t(\tau_i(x)) - \sum_{i=1}^{\tilde{n}} \tilde{w}_i \nabla_{\theta} g'_t(\tau_i(x)) \right\|_1 \\
&\quad + \left\| \frac{\nabla_{\theta} g'_t(x)}{g'_t(x)} - \frac{\nabla_{\theta} f'_t(x)}{f'_t(x)} \right\|_1 \\
&= \text{I} + \text{II}
\end{aligned}$$

where

$$\begin{aligned}
\text{I} &= \left\| \sum_{i=1}^{\tilde{n}} \tilde{w}_i \nabla_{\theta} f'_t(\tau_i(x)) - \sum_{i=1}^{\tilde{n}} \tilde{w}_i \nabla_{\theta} g'_t(\tau_i(x)) \right\|_1 \\
\text{II} &= \left\| \frac{\nabla_{\theta} g'_t(x)}{g'_t(x)} - \frac{\nabla_{\theta} f'_t(x)}{f'_t(x)} \right\|_1
\end{aligned}$$

First proving bound on I,

$$\begin{aligned}
\text{I} &= \left\| \sum_{i=1}^{\tilde{n}} \tilde{w}_i \nabla_{\theta} f'_t(\tau_i(x)) - \sum_{i=1}^{\tilde{n}} \tilde{w}_i \nabla_{\theta} g'_t(\tau_i(x)) \right\|_1 \\
&\leq \sum_{i=1}^{\tilde{n}} \tilde{w}_i \left\| \nabla_{\theta} f'_t(\tau_i(x)) - \nabla_{\theta} g'_t(\tau_i(x)) \right\|_1 \\
&\stackrel{(i)}{\leq} 8c_1 \epsilon_a m \Delta_{np}^t \sqrt{2 \log m} + \left(\frac{128 \eta m^{1.5} \bar{\Delta} c_1 c_2 \epsilon_a t \sqrt{\log m}}{\sqrt{\pi}} \right)
\end{aligned}$$

where inequality (i) follows from lemma 7. Now, we will bound II,

$$\begin{aligned}
\text{II} &= \left\| \frac{\nabla_{\theta} g'_t(x)}{g'_t(x)} - \frac{\nabla_{\theta} f'_t(x)}{f'_t(x)} \right\|_1 \\
&= \left\| \frac{\exp(P_t(x)) \mathbb{I}[P_t(x) < 0] + \mathbb{I}[P_t(x) \geq 0]}{\exp(P_t(x)) \mathbb{I}[P_t(x) < 0] + (P_t(x) + 1) \mathbb{I}[P_t(x) \geq 0]} \nabla_{\theta} P_t(x) \right. \\
&\quad \left. - \frac{\exp(N_t(x)) \mathbb{I}[N_t(x) < 0] + \mathbb{I}[N_t(x) \geq 0]}{\exp(N_t(x)) \mathbb{I}[N_t(x) < 0] + (N_t(x) + 1) \mathbb{I}[N_t(x) \geq 0]} \nabla_{\theta} N_t(x) \right\|_1 \\
&= \left\| \left(\mathbb{I}[P_t(x) < 0] + \frac{\mathbb{I}[P_t(x) \geq 0]}{(P_t(x) + 1)} \right) \nabla_{\theta} P_t(x) - \left(\mathbb{I}[N_t(x) < 0] + \frac{\mathbb{I}[N_t(x) \geq 0]}{(N_t(x) + 1)} \right) \nabla_{\theta} N_t(x) \right\|_1 \\
&= \text{II}_1 + \text{II}_2 + \text{II}_3 + \text{II}_4
\end{aligned}$$

where

$$\begin{aligned}\Pi_1 &= \left\| \nabla_{\theta} P_t(x) - \nabla_{\theta} N_t(x) \right\|_1 \mathbb{I}[P_t(x) < 0, N_t(x) < 0] \\ \Pi_2 &= \left\| \nabla_{\theta} P_t(x) - \frac{\nabla_{\theta} N_t(x)}{N_t(x) + 1} \right\|_1 \mathbb{I}[P_t(x) < 0, N_t(x) \geq 0] \\ \Pi_3 &= \left\| \frac{\nabla_{\theta} P_t(x)}{P_t(x) + 1} - \nabla_{\theta} N_t(x) \right\|_1 \mathbb{I}[P_t(x) \geq 0, N_t(x) < 0] \\ \Pi_4 &= \left\| \frac{\nabla_{\theta} P_t(x)}{P_t(x) + 1} - \frac{\nabla_{\theta} N_t(x)}{N_t(x) + 1} \right\|_1 \mathbb{I}[P_t(x) \geq 0, N_t(x) \geq 0]\end{aligned}$$

Bounding Π_2 , we get

$$\begin{aligned}\Pi_2 &\leq \left(\left| \frac{1}{N_t(x) + 1} \right| \left\| \nabla_{\theta} P_t(x) - \nabla_{\theta} N_t(x) \right\|_1 + \left| \frac{N_t(x)}{1 + N_t(x)} \right| \left\| \nabla_{\theta} P_t(x) \right\|_1 \right) \mathbb{I}[P_t(x) < 0, N_t(x) \geq 0] \\ &\leq \left(\left\| \nabla_{\theta} P_t(x) - \nabla_{\theta} N_t(x) \right\|_1 + \Delta_{np}^t \left\| \nabla_{\theta} P_t(x) \right\|_1 \right) \mathbb{I}[P_t(x) < 0, N_t(x) \geq 0]\end{aligned}\quad (23)$$

Similarly, bounding Π_3 , we get

$$\begin{aligned}\Pi_3 &\leq \left(\left| \frac{1}{P_t(x) + 1} \right| \left\| \nabla_{\theta} P_t(x) - \nabla_{\theta} N_t(x) \right\|_1 + \left| \frac{P_t(x)}{1 + P_t(x)} \right| \left\| \nabla_{\theta} N_t(x) \right\|_1 \right) \mathbb{I}[P_t(x) \geq 0, N_t(x) < 0] \\ &\leq \left(\left\| \nabla_{\theta} P_t(x) - \nabla_{\theta} N_t(x) \right\|_1 + \Delta_{np}^t \left\| \nabla_{\theta} N_t(x) \right\|_1 \right) \mathbb{I}[P_t(x) \geq 0, N_t(x) < 0]\end{aligned}\quad (24)$$

Bounding Π_4 , we get

$$\begin{aligned}\Pi_4 &\leq \left(\left\| \frac{\nabla_{\theta} P_t(x)}{P_t(x) + 1} - \frac{\nabla_{\theta} N_t(x)}{P_t(x) + 1} \right\|_1 + \left\| \frac{\nabla_{\theta} N_t(x)}{P_t(x) + 1} - \frac{\nabla_{\theta} N_t(x)}{N_t(x) + 1} \right\|_1 \right) \mathbb{I}[P_t(x) \geq 0, N_t(x) \geq 0] \\ &\leq \left(\frac{1}{P_t(x) + 1} \left\| \nabla_{\theta} P_t(x) - \nabla_{\theta} N_t(x) \right\|_1 + \frac{\left\| \nabla_{\theta} N_t(x) \right\|_1 \Delta_{np}^t}{(P_t(x) + 1)(N_t(x) + 1)} \right) \mathbb{I}[P_t(x) \geq 0, N_t(x) \geq 0] \\ &\leq \left(\left\| \nabla_{\theta} P_t(x) - \nabla_{\theta} N_t(x) \right\|_1 + \left\| \nabla_{\theta} N_t(x) \right\|_1 \Delta_{np}^t \right) \mathbb{I}[P_t(x) \geq 0, N_t(x) \geq 0]\end{aligned}\quad (25)$$

Using eq.(23), eq.(24) and eq.(25), we get

$$\begin{aligned}\Pi &= \left\| \frac{\nabla_{\theta} g'_t(x)}{g'_t(x)} - \frac{\nabla_{\theta} f'_t(x)}{f'_t(x)} \right\|_1 \\ &\leq \left\| \nabla_{\theta} P_t(x) - \nabla_{\theta} N_t(x) \right\|_1 + \left\| \nabla_{\theta} N_t(x) \right\|_1 \Delta_{np}^t \mathbb{I}[P_t(x) \geq 0] \\ &\quad + \Delta_{np}^t \left\| \nabla_{\theta} P_t(x) \right\|_1 \mathbb{I}[P_t(x) < 0, N_t(x) \geq 0]\end{aligned}$$

Using eq.(21) and eq.(22), we get

$$\begin{aligned}\Pi &\leq \left(\frac{64\eta m^{1.5} \bar{\Delta} c_1 c_2 \epsilon_a t \sqrt{\log m}}{\sqrt{\pi}} \right) + \Delta_{np}^t \left(\left\| \nabla_{\theta} N_t(x) \right\|_1 + \left\| \nabla_{\theta} P_t(x) \right\|_1 \right) \\ &\leq \left(\frac{64\eta m^{1.5} \bar{\Delta} c_1 c_2 \epsilon_a t \sqrt{\log m}}{\sqrt{\pi}} \right) + \Delta_{np}^t \left(\left\| (A_0, A_0) \odot (x, 1) \odot (\mathbb{I}[w_{r0}x + b_{r0}], \mathbb{I}[w_{r0}x + b_{r0}]) \right\|_1 \right. \\ &\quad \left. + \left\| (A_0, A_0) \odot (x, 1) \odot (\mathbb{I}[(w_{r0} + w_r^t)x + b_{r0} + b_r^t], \mathbb{I}[(w_{r0} + w_r^t)x + b_{r0} + b_r^t]) \right\|_1 \right) \\ &\leq \left(\frac{64\eta m^{1.5} \bar{\Delta} c_1 c_2 \epsilon_a t \sqrt{\log m}}{\sqrt{\pi}} \right) + \Delta_{np}^t \left(8c_1 \epsilon_a m \sqrt{2 \log m} \right) \\ &\leq \left(\frac{64\eta m^{1.5} \bar{\Delta} c_1 c_2 \epsilon_a t \sqrt{\log m}}{\sqrt{\pi}} \right) + 8c_1 \epsilon_a m \Delta_{np}^t \sqrt{2 \log m}\end{aligned}\quad (26)$$

Combining bounds on I and II, we get

$$\begin{aligned} \|\nabla_{\theta}\hat{L}(f'_t, x) - \nabla_{\theta}\hat{L}(g'_t, x)\|_1 &\leq 8c_1\epsilon_a m\Delta_{np}^t\sqrt{2\log m} + \frac{128\eta m^{1.5}\bar{\Delta}c_1c_2\epsilon_a t\sqrt{\log m}}{\sqrt{\pi}} \\ &\quad + \frac{64\eta m^{1.5}\bar{\Delta}c_1c_2\epsilon_a t\sqrt{\log m}}{\sqrt{\pi}} + 8c_1\epsilon_a m\Delta_{np}^t\sqrt{2\log m} \\ &\leq \frac{192\eta m^{1.5}\bar{\Delta}c_1c_2\epsilon_a t\sqrt{\log m}}{\sqrt{\pi}} + 16c_1\epsilon_a m\Delta_{np}^t\sqrt{2\log m} \end{aligned}$$

Define Δ as the upper bound on $\|\nabla_{\theta}\hat{L}(f'_t, x) - \nabla_{\theta}\hat{L}(g'_t, x)\|_1$.

$$\Delta = \frac{192\eta m^{1.5}\bar{\Delta}c_1c_2\epsilon_a t\sqrt{\log m}}{\sqrt{\pi}} + 16c_1\epsilon_a m\Delta_{np}^t\sqrt{2\log m} \quad (27)$$

□

C OPTIMIZATION PROOF

This section shows that gradient-based optimization of the loss for the target function can be closely approximated by the gradient-based optimization of the pseudo network. Since the loss function of the pseudo network is convex in its parameters, we get global optimization.

Lemma 9. (*Convexity of loss function of pseudo network*) *The loss function for pseudo network is convex with respect to parameters of neural network.*

Proof. The loss function for pseudo network is

$$\hat{L}(g'_t, x) = \sum_{i=1}^{\tilde{n}} \tilde{w}_i g'_t(\tau_i(x)) - \log(g'_t(x))$$

Dividing the loss function in 2 parts,

$$\hat{L}(g'_t, x) = \hat{L}_1(g'_t, x) + \hat{L}_2(g'_t, x)$$

where

$$\hat{L}_1(g'_t, x) = \sum_{i=1}^{\tilde{n}} \tilde{w}_i g'_t(\tau_i(x))$$

$$\hat{L}_2(g'_t, x) = -\log(g'_t(x))$$

We will prove convexity of both $\hat{L}_1(g'_t, x)$ and $\hat{L}_2(g'_t, x)$. To prove convexity of $\hat{L}_1(g'_t, x)$ as a function of parameters θ , we will prove that Hessian of $\hat{L}_1(g'_t, x)$ is positive semidefinite.

$$\begin{aligned} \nabla_{\theta}\hat{L}_1(g'_t, x) &= \sum_{i=1}^{\tilde{n}} \tilde{w}_i \nabla_{\theta} g'_t(\tau_i(x)) = \sum_{i=1}^{\tilde{n}} \tilde{w}_i \phi'(P_t(\tau_i(x))) \nabla_{\theta} P_t(\tau_i(x)) \\ \nabla_{\theta}^2 \hat{L}_1(g'_t, x) &= \sum_{i=1}^{\tilde{n}} \tilde{w}_i \nabla_{\theta}^2 g'_t(\tau_i(x)) \\ &= \sum_{i=1}^{\tilde{n}} \tilde{w}_i \phi''(P_t(\tau_i(x))) \nabla_{\theta} P_t(\tau_i(x)) \nabla_{\theta} P_t(\tau_i(x))^T \\ &\quad + \sum_{i=1}^{\tilde{n}} \tilde{w}_i \phi'(P_t(\tau_i(x))) \nabla_{\theta}^2 P_t(\tau_i(x)) \\ &= \sum_{i=1}^{\tilde{n}} \tilde{w}_i \phi''(P_t(\tau_i(x))) \nabla_{\theta} P_t(\tau_i(x)) \nabla_{\theta} P_t(\tau_i(x))^T \end{aligned}$$

The first term of the Hessian matrix is sum of Gram matrix and the second term of the Hessian matrix is Gram matrix. Hence, the Hessian of $\hat{L}_1(g'_t, x)$ is positive semidefinite. For second term,

$$\begin{aligned}\hat{L}_2(g'_t, x) &= -\log(\exp(P_t(x))\mathbb{I}[P_t(x) \leq 0] + (P_t(x) + 1)\mathbb{I}[P_t(x) > 0]) \\ &= -P_t(x)\mathbb{I}[P_t(x) \leq 0] - \log(P_t(x) + 1)\mathbb{I}[P_t(x) > 0]\end{aligned}$$

Note that $\hat{L}_2(g'_t, x)$ is convex in $P_t(x)$ and $P_t(x)$ is linear in θ . Composition of convex and linear function is convex therefore, $\hat{L}_2(g'_t, x)$ is convex in θ . As sum of 2 convex functions is convex, $\hat{L}(g'_t, x)$ is convex. \square

Remark 1. *If we use base distribution as standard Gaussian distribution, then loss function will have following term.*

$$\left(\sum_{i=1}^{\tilde{n}} \tilde{w}_i \nabla_{\theta} g'_t(\tau_i(x))\right) \sum_{i=1}^{\tilde{n}} \tilde{w}_i \phi''(P_t(\tau_i(x))) \nabla_{\theta} P_t(\tau_i(x)) \nabla_{\theta} P_t(\tau_i(x))^T$$

The above term won't remain convex in parameters of neural network θ if we use standard Gaussian distribution as base distribution. If base distribution is standard Gaussian distribution, then \tilde{g} has to be negative for some points. At points with negative values of \tilde{g} , Hessian of $\hat{L}_1(g'_t, x)$ with respect to θ will be negative semidefinite and $\hat{L}_1(g'_t, x)$ will not remain convex in θ .

Lemma 10. (Approximated loss is close to optimal loss) *For every $\epsilon \in (0, 1)$, there exist $m > \text{poly}(U_h, \frac{1}{\epsilon})$, $\eta = \tilde{O}(\frac{1}{m\epsilon})$ and $T = O\left(\frac{U_h^2 \log m}{\epsilon^2}\right)$ such that, with at least $0.93 - 2 \exp\left(-\frac{m}{8\epsilon_a^2}\right) - \exp\left(-\frac{\epsilon^2 m}{128c_1^2 U_h^2 \log m}\right)$ probability, we get*

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_{\text{sgd}}[\hat{L}(f'_t, \mathcal{X})] - \hat{L}(F^{*'}, \mathcal{X}) \leq O(\epsilon')$$

Proof. For set of examples \mathcal{X} , define

$$\hat{L}(f'_t, \mathcal{X}) = \frac{1}{|\mathcal{X}|} \sum_{x \in \mathcal{X}} \hat{L}(f'_t, x)$$

Here, $|\mathcal{X}|$ denotes size of set \mathcal{X} . From lemma 9, we know that $\hat{L}(g'_t, \mathcal{X})$ is convex in parameters θ . Using convexity of $\hat{L}(g'_t, \mathcal{X})$ wrt θ ,

$$\begin{aligned}\hat{L}(g'_t, \mathcal{X}) - \hat{L}(g^*, \mathcal{X}) &\leq \langle \nabla_{\theta} \hat{L}(g'_t, \mathcal{X}), \theta^t - \theta^* \rangle \\ &\leq \|\nabla_{\theta} \hat{L}(g'_t, \mathcal{X}) - \nabla_{\theta} \hat{L}(f'_t, \mathcal{X})\|_1 \|\theta^t - \theta^*\|_{\infty} \\ &\quad + \langle \nabla_{\theta} \hat{L}(f'_t, \mathcal{X}), \theta^t - \theta^* \rangle\end{aligned}\tag{28}$$

where $\|\cdot\|_1$ and $\|\cdot\|_{\infty}$ denotes l_1 and l_{∞} norm respectively. x_t is a training point. The stochastic gradient descent updates the parameters using x_t at time t .

$$\begin{aligned}\|\theta^{t+1} - \theta^*\|_2^2 &= \|\tilde{\theta}^{t+1} - \theta^*\|_2^2 \\ &= \|\theta^t - \eta \nabla_{\theta} \hat{L}(f'_t, x^t) - \theta^*\|_2^2 \\ &= \|\theta^t - \theta^*\|_2^2 + \eta^2 \|\nabla_{\theta} \hat{L}(f'_t, x^t)\|_2^2 - 2\eta \langle \theta^t - \theta^*, \nabla_{\theta} \hat{L}(f'_t, x^t) \rangle\end{aligned}$$

Taking expectation wrt x_t ,

$$\mathbb{E}_{x^t} [\|\theta^{t+1} - \theta^*\|_2^2] = \|\theta^t - \theta^*\|_2^2 + \eta^2 \mathbb{E}_{x^t} [\|\nabla_{\theta} \hat{L}(f'_t, x^t)\|_2^2] - 2\eta \langle \nabla_{\theta} \hat{L}(f'_t, \mathcal{X}), \theta^t - \theta^* \rangle\tag{29}$$

Using eq.(29) and eq.(28),

$$\begin{aligned}\hat{L}(g'_t, \mathcal{X}) - \hat{L}(g^*, \mathcal{X}) &\leq \|\nabla_{\theta} \hat{L}(g'_t, \mathcal{X}) - \nabla_{\theta} \hat{L}(f'_t, x)\|_1 \|\theta^t - \theta^*\|_{\infty} \\ &\quad + \frac{\|\theta^t - \theta^*\|_2^2 - \mathbb{E}_{x^t} [\|\theta^{t+1} - \theta^*\|_2^2]}{2\eta} \\ &\quad + \frac{\eta}{2} \mathbb{E}_{x^t} [\|\nabla_{\theta} \hat{L}(f'_t, x^t)\|_2^2]\end{aligned}$$

Using eq.(16) and eq.(17), we get

$$\left\| \nabla_{\theta} \hat{L}(f'_t, x^t) \right\|_2^2 \leq 2m\bar{\Delta}^2$$

Averaging from $t = 0$ to $T - 1$, Using eq.(1) we get

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_{sgd}[\hat{L}(g'_t, \mathcal{X})] - \hat{L}(g^*, \mathcal{X}) \leq \Delta \left(\sup_{t \in [T]} \|\theta_t\|_{\infty} + \|\theta^*\|_{\infty} \right) + \frac{\|\theta^*\|_2^2}{2\eta T} + \eta m \bar{\Delta}^2$$

Using lemma 3 and lemma 6, with at least $0.99 - \frac{1}{c_1} - \sum_{t=1}^T \exp\left(-\frac{64(1-c_2)^2 \eta^2 m^2 \bar{\Delta}^2 t^2}{\pi}\right) - \frac{1}{c_3} - \frac{1}{c_6} - \frac{1}{c_7} - 2 \exp\left(-\frac{m}{8\epsilon_a^2}\right) - \exp\left(-\frac{\epsilon^2 m}{128c_1^2 U_h^2 \log m}\right)$ probability, we get

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_{sgd}[\hat{L}(f'_t, \mathcal{X})] - \hat{L}(F^{*'}, \mathcal{X}) &\leq \Delta \left(\sup_{t \in [T]} \|\theta_t\|_{\infty} + \|\theta^*\|_{\infty} \right) + \frac{\|\theta^*\|_2^2}{2\eta T} + \eta m \bar{\Delta}^2 \\ &\quad + 3\Delta_{np}^t + 3 \left(16c_1 (c_6 + c_7) \epsilon_a \log m + \omega_{\phi^{-1}(F^{*'})}(\delta) + \epsilon \right) \end{aligned}$$

We choose following values/relations of η, T and $m_{F^{*'}}$

$$\begin{aligned} \eta &= \frac{\epsilon}{m\bar{\Delta}^2} \\ &= \frac{\epsilon}{m(6c_1 \epsilon_a \sqrt{2 \log m})^2} \\ &= \frac{\epsilon}{72c_1^2 m \epsilon_a^2 \log m} \\ T &= \frac{\|\theta^*\|_2^2}{2\eta\epsilon} \\ &= \frac{U_h^2 \pi}{2m\epsilon_a^2} \frac{72c_1^2 m \epsilon_a^2 \log m}{2\epsilon^2} \\ &= \frac{18\pi c_1^2 \log m U_h^2}{\epsilon^2} \end{aligned} \tag{30}$$

We can choose δ such that $\omega_{\phi^{-1}(F^{*'})}(\delta) = \epsilon$. Using above inequalities, we get following equalities.

$$\begin{aligned} \frac{\|\theta^*\|_2^2}{2\eta T} &= \frac{\|\theta^*\|_2^2}{2\eta} \frac{2\eta\epsilon}{\|\theta^*\|_2^2} = \epsilon \\ \eta m \bar{\Delta}^2 &= \frac{\epsilon}{m\bar{\Delta}^2} m \bar{\Delta}^2 = \epsilon \end{aligned}$$

$$3 \left(16c_1 (c_6 + c_7) \epsilon_a \log m + \omega_{\phi^{-1}(F^{*'})}(\delta) + \epsilon \right) \leq 3 \left(16c_1 (c_6 + c_7) \epsilon_a \log m + 2\epsilon \right)$$

To get value of m ,

$$\begin{aligned}
\sup_{t \in [T]} \|\theta^t\|_\infty &= \sup_{t \in [T]} \eta \bar{\Delta} t = \eta \bar{\Delta} T = \frac{\|\theta^*\|_2^2 \bar{\Delta}}{2\epsilon} = \frac{U_h^2 \pi}{2m\epsilon_a^2} \frac{(6c_1 \epsilon_a \sqrt{2 \log m})}{2\epsilon} \\
&\leq \frac{3\pi U_h^2 c_1 \sqrt{\log m}}{\sqrt{2} m \epsilon_a \epsilon} \\
\|\theta^*\|_\infty &= \frac{U_h \sqrt{\pi}}{\sqrt{2} m \epsilon_a} \\
\sup_{t \in [T]} \|\theta^t\|_\infty + \|\theta^*\|_\infty &\leq \frac{\pi U_h^2 (1 + 3c_1) \sqrt{\log m}}{\sqrt{2} m \epsilon_a \epsilon} \\
\Delta &= \frac{192\eta m^{1.5} \bar{\Delta} c_1 c_2 \epsilon_a t \sqrt{\log m}}{\sqrt{\pi}} + 16c_1 \epsilon_a m \Delta_{np}^t \sqrt{2 \log m} \\
&\leq \frac{192\eta m^{1.5} \bar{\Delta} c_1 c_2 \epsilon_a t \sqrt{\log m}}{\sqrt{\pi}} + 16c_1 \epsilon_a m \sqrt{2 \log m} \left(\frac{192\eta^2 m^{1.5} \bar{\Delta}^2 c_1 c_2 \epsilon_a t^2 \sqrt{\log m}}{\sqrt{\pi}} \right) \\
&\leq \frac{192\eta m^{1.5} \bar{\Delta} c_1 c_2 \epsilon_a t \sqrt{\log m}}{\sqrt{\pi}} + \frac{3072\sqrt{2} c_1^2 c_2 \epsilon_a^2 \eta^2 t^2 m^{2.5} \log m \bar{\Delta}^2}{\sqrt{\pi}} \\
&\leq \frac{192m^{1.5} c_1 c_2 \epsilon_a \sqrt{\log m}}{\sqrt{\pi}} \left(\frac{U_h^2 \pi}{4m\epsilon_a^2 \epsilon} \right) (6c_1 \epsilon_a \sqrt{2 \log m}) \\
&\quad + \frac{3072\sqrt{2} c_1^2 c_2 \epsilon_a^2 m^{2.5} \log m}{\sqrt{\pi}} \left(\frac{U_h^2 \pi}{4m\epsilon_a^2 \epsilon} \right)^2 (6c_1 \epsilon_a \sqrt{2 \log m})^2 \\
&\leq \frac{288\sqrt{2\pi} \sqrt{m} \log m c_1^2 c_2 U_h^2}{\epsilon} + \frac{13824\sqrt{2\pi^3} c_1^4 c_2 \sqrt{m} (\log m)^2 U_h^4}{\epsilon^2} \\
&\leq \frac{14112\sqrt{2\pi^3} c_1^4 c_2 \sqrt{m} (\log m)^2 U_h^4}{\epsilon^2}
\end{aligned}$$

Multiplication of Δ and $\left(\sup_{t \in [T]} \|\theta^t\|_\infty + \|\theta^*\|_\infty \right)$ will be

$$\begin{aligned}
\Delta \left(\sup_{t \in [T]} \|\theta^t\|_\infty + \|\theta^*\|_\infty \right) &\leq \frac{14112\sqrt{2\pi^3} c_1^4 c_2 \sqrt{m} (\log m)^2 U_h^4}{\epsilon^2} \left(\frac{\pi U_h^2 (1 + 3c_1) \sqrt{\log m}}{\sqrt{2} m \epsilon_a \epsilon} \right) \\
&= \frac{14112\pi^{2.5} c_1^4 c_2 (1 + 3c_1) (\log m)^{2.5} U_h^6}{\sqrt{m} \epsilon_a \epsilon^3}
\end{aligned}$$

Taking m as

$$m \geq \Omega \left(\frac{c_1^8 c_2^2 (1 + 3c_1)^2 U_h^{12}}{\epsilon_a^2 \epsilon^8} \right) \quad (31)$$

Choosing m which satisfies above inequality will give us the following inequality.

$$\Delta \left(\sup_{t \in [T]} \|\theta^t\|_\infty + \|\theta^*\|_\infty \right) \leq \epsilon$$

Using eq.(21), we get

$$\begin{aligned}
\Delta_{np}^t &= \left(\frac{192\eta^2 m^{1.5} \bar{\Delta}^2 c_1 c_2 \epsilon_a t^2 \sqrt{\log m}}{\sqrt{\pi}} \right) \\
&\leq \left(\frac{192m^{1.5} c_1 c_2 \epsilon_a \sqrt{\log m}}{\sqrt{\pi}} \right) \left(\frac{U_h^2 \pi}{4m\epsilon_a^2 \epsilon} \right)^2 (6c_1 \epsilon_a \sqrt{2 \log m})^2 \\
&\leq \left(\frac{864\pi^{1.5} c_1^3 c_2 (\log m)^{1.5} U_h^4}{\sqrt{m} \epsilon^2 \epsilon_a} \right)
\end{aligned} \quad (32)$$

(33)

Using the choice of m , we get that $\Delta_{np}^t \leq \epsilon$. Using eq.(30) and eq.(31), with at least $0.99 - \frac{1}{c_1} - \frac{1}{c_3} - \frac{1}{c_6} - \frac{1}{c_7} - \sum_{t=1}^T \exp\left(-\frac{64(1-c_2)^2 \eta^2 m^2 \bar{\Delta}^2 t^2}{\pi}\right) - \exp\left(-\frac{\epsilon^2 m}{128c_1^2 U_h^2 \log m}\right)$ probability, we get

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_{sgd}[\hat{L}(f'_t, \mathcal{X})] - \hat{L}(F^{*'}, \mathcal{X}) &\leq \Delta \left(\sup_{t \in [T]} \|\theta_t\|_\infty + \|\theta^*\|_\infty \right) + \frac{\|\theta^*\|_2^2}{2\eta T} + \eta m \bar{\Delta}^2 \\ &\quad + \frac{2592\pi^{1.5} c_1^3 c_2 (\log m)^{1.5} U_h^4}{\sqrt{m\epsilon^2 \epsilon_a}} + 3(16c_1(c_6 + c_7)\epsilon_a \log m + 2\epsilon) \\ &\leq 3\epsilon + \frac{2592\pi^{1.5} c_1^3 c_2 (\log m)^{1.5} U_h^4}{\sqrt{m\epsilon^2 \epsilon_a}} + 3(16c_1(c_6 + c_7)\epsilon_a \log m + 2\epsilon) \end{aligned}$$

Taking $c_1 = 100, c_2 = 0.01, c_3 = 100, c_5 = 1000, c_6 = 100, c_7 = 100, \epsilon_a = \frac{\epsilon}{6000 \log m} \leq \epsilon$, with at least $0.95 - \sum_{t=1}^T \exp\left(-\frac{64(0.9801)\eta^2 m^2 \bar{\Delta}^2 t^2}{\pi}\right) - 2 \exp\left(-\frac{m}{8\epsilon_a^2}\right) - \exp\left(-\frac{\epsilon^2 m}{128c_1^2 U_h^2 \log m}\right)$

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_{sgd}[\hat{L}(f'_t, \mathcal{X})] - \hat{L}(F^{*'}, \mathcal{X}) &\leq 3\epsilon + \frac{2592\pi^{1.5} c_1^3 c_2 (\log m)^{1.5} U_h^4}{\sqrt{m\epsilon^2 \epsilon_a}} + O(\epsilon) \\ &\leq 3\epsilon + O\left(\frac{\epsilon^2}{U_h^2}\right) + O(\epsilon) \end{aligned}$$

For any $\epsilon \in [0, 1]$, with probability at least $0.95 - \sum_{t=1}^T \exp\left(-\frac{64(0.9801)\eta^2 m^2 \bar{\Delta}^2 t^2}{\pi}\right) - 2 \exp\left(-\frac{m}{8\epsilon_a^2}\right) - \exp\left(-\frac{\epsilon^2 m}{128c_1^2 U_h^2 \log m}\right)$

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_{sgd}[\hat{L}(f'_t, \mathcal{X})] - \hat{L}(F^{*'}, \mathcal{X}) \leq O(\epsilon)$$

To find lower bound on probability, we use $\sum_{t=1}^T \frac{1}{t^2} \leq \sum_{t=1}^\infty \frac{1}{t^2} \leq 2$.

$$\begin{aligned} \sum_{t=1}^T \exp\left(-\frac{64(0.9801)\eta^2 m^2 \bar{\Delta}^2 t^2}{\pi}\right) &\leq \sum_{t=1}^T \frac{\pi}{64(0.9801)\eta^2 m^2 \bar{\Delta}^2 t^2} \\ &\leq \frac{\pi (m\bar{\Delta}^2)^2}{32(0.9801)\epsilon^2 m^2 \bar{\Delta}^2} \\ &\leq \frac{\pi \bar{\Delta}^2}{32(0.9801)\epsilon^2} \\ &\leq \frac{\pi}{3200(0.9801)} \\ &\leq 0.01 \end{aligned}$$

Finally, with at least $0.93 - 2 \exp\left(-\frac{m}{8\epsilon_a^2}\right) - \exp\left(-\frac{\epsilon^2 m}{128c_1^2 U_h^2 \log m}\right)$ probability,

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_{sgd}[\hat{L}(f'_t, \mathcal{X})] - \hat{L}(F^{*'}, \mathcal{X}) \leq O(\epsilon)$$

□

D GENERALIZATION

In this section, we prove generalization guarantees to complement our optimization result, and complete the proof of our main theorem (Theorem 2) about efficiently learning distributions using

univariate normalizing flows. Recall that the approximate loss function \hat{L} is given by

$$\begin{aligned}\hat{L}(f'_t, x) &= \hat{L}(N_t(\tau_1(x)), N_t(\tau_2(x)), \dots, N_t(\tau_{\tilde{n}}(x)), N_t(x)) \\ &= \sum_{i=1}^{\tilde{n}} \tilde{w}_i \phi(N_t(\tau_i(x))) - \log(\phi(N_t(x)))\end{aligned}$$

where

$$N(x) = \sum_{r=1}^m a_{r0} \sigma((w_{r0} + w_r)x + (b_{r0} + b_r))$$

Lemma 11. (Empirical Rademacher complexity for two-layer neural network) For every $B > 0$, for every $n \geq 1$, with at least $1 - \frac{1}{c_1}$ probability, the empirical Rademacher complexity is bounded by

$$\frac{1}{n} \mathbb{E}_{\xi \in \{\pm 1\}^n} \left[\sup_{\max_{r \in [m]} |w_r|, |b_r| \leq B} \sum_{i=1}^n \xi_i N(x_i) \right] \leq \frac{8c_1 \epsilon_a B m \sqrt{2 \log m}}{\sqrt{n}}$$

Proof. Using part a of proposition 1, we get that $\{x \rightarrow w_r x + b_r \mid |w_r| \leq B, |b_r| \leq B\}$ has Rademacher complexity $\frac{2B}{\sqrt{n}}$. Using part b of lemma 1, we get that $\{x \rightarrow ((w_{r0} + w_r)x + (b_{r0} + b_r)) \mid |w_r| \leq B, |b_r| \leq B, w_{r0}, b_{r0} \sim \mathcal{N}(0, \frac{1}{m})\}$ has Rademacher complexity $\frac{2B}{\sqrt{n}}$. Using part c of lemma 1, we get that class of functions in $\mathcal{F} = \{x \rightarrow N(x) \mid \max_{r \in [m]} |w_r| \leq B, \max_{r \in [m]} |b_r| \leq B\}$ has Rademacher complexity

$$\hat{\mathcal{R}}(\mathcal{X}; \mathcal{F}) \leq 2 \|\mathbf{a}\|_1 \frac{2B}{\sqrt{n}} \stackrel{(i)}{\leq} \frac{8c_1 \epsilon_a B m \sqrt{2 \log m}}{\sqrt{n}}$$

where inequality (i) follows from lemma 15 with at least $1 - \frac{1}{c_1}$ probability. \square

Lemma 12. Suppose n is sufficiently high such that it satisfies following condition.

$$n \geq O\left(\frac{(M_{\hat{L}} - m_{\hat{L}})^2 (Q+1)^2 U_h^4 (\log m)^2}{\epsilon^2}\right)$$

If n satisfies above condition, then with at least 0.98 probability, population loss of any functions of set $\{x \rightarrow N(x) \mid |w_r| \leq \eta \bar{\Delta} T, |b_r| \leq \eta \bar{\Delta} T \forall r \in [m]\}$ is close to empirical loss i.e.

$$\sup_{N \in \mathcal{F}} \left| \mathbb{E}_{x \in \mathcal{D}} \left[\hat{L}(N(\tau_1(x)), \dots, N(\tau_Q(x)), N(x)) \right] - \frac{1}{n} \sum_{i=1}^n \hat{L}(N(\tau_1(x_i)), \dots, N(\tau_Q(x_i)), N(x_i)) \right| \leq \epsilon$$

Proof. Using Fact 7, with at least $1 - \delta$ probability, we get

$$\begin{aligned}\sup_{N \in \mathcal{F}} \left| \mathbb{E}_{x \sim \mathcal{D}} \left[\hat{L}(N_t(\tau_1(x)), \dots, N_t(\tau_{\tilde{n}}(x)), N_t(x)) \right] - \frac{1}{n} \sum_{i=1}^n \hat{L}(N(\tau_1(x_i)), \dots, N(\tau_{\tilde{n}}(x_i)), N(x_i)) \right| \\ \leq 2\sqrt{2} L_s (\tilde{n} + 1) \hat{\mathcal{R}}(\mathcal{X}; \mathcal{F}) + b \sqrt{\frac{\log \frac{1}{\delta}}{2n}}\end{aligned} \tag{34}$$

where $\mathcal{F} = \{x \rightarrow N(x) \mid |w_r| \leq \eta\bar{\Delta}T, |b_r| \leq \eta\bar{\Delta}T \ \forall r \in [m]\}$. We get coordinate wise Lipschitz continuity of loss \hat{L} function as following.

$$\begin{aligned} L_j &\leq \sup_{N \in \mathcal{F}, |x| \leq 1} |\tilde{w}_j \phi'(N(\tau_j(x)))| \\ &\leq \sup_{N \in \mathcal{F}, |x| \leq 1} \frac{1}{\tilde{n}} |\phi'(N(\tau_j(x)))| \\ &\leq \frac{1}{\tilde{n}} \quad \forall i \in [\tilde{n}] \\ L_{\tilde{n}+1} &\leq \sup_{N \in \mathcal{F}, |x| \leq 1} \frac{\phi'(N(x))}{\phi(N(x))} = \sup_{N \in \mathcal{F}, |x| \leq 1} \frac{\exp(N(x)) \mathbb{I}[N(x) \leq 0] + \mathbb{I}[N(x) \geq 0]}{\exp(N(x)) \mathbb{I}[N(x) \leq 0] + (N(x) + 1) \mathbb{I}[N(x) \geq 0]} \\ &\leq \sup_{N \in \mathcal{F}, |x| \leq 1} \mathbb{I}[N(x) \leq 0] + \frac{1}{N(x) + 1} \mathbb{I}[N(x) \geq 0] \\ &\leq 1 \end{aligned}$$

Using lemma 17, standard Lipschitz constant of \hat{L} is

$$L_s \leq \sqrt{\sum_{i=1}^{Q+1} L_i^2} \leq \sqrt{\frac{1}{Q} + 1} \leq 2 \quad (35)$$

To get upper bound on \hat{L} , we use Lipschitz property of \hat{L} .

$$\left| \hat{L}(N(\tau_1(x)), N(\tau_2(x)), \dots, N(\tau_{\tilde{n}}(x)), N(x)) - \hat{L}(0, 0, \dots, 0, 0) \right| \leq \sum_{i=1}^{\tilde{n}} \frac{1}{\tilde{n}} |N(\tau_i(x))| + |N(x)| \quad (36)$$

Finding upper bound $N(x)$ for all $x \in [-1, 1]$,

$$\begin{aligned} \sup_{N \in \mathcal{F}, x \in [-1, 1]} N(x) &\leq \sup_{|w_r| \leq \eta\bar{\Delta}T, |b_r| \leq \eta\bar{\Delta}T, x \in [-1, 1]} P_t(x) + \Delta_{np}^T \\ &\leq \sup_{|w_r| \leq \eta\bar{\Delta}T, |b_r| \leq \eta\bar{\Delta}T, x \in [-1, 1]} \sum_{r=1}^m a_{r0} \sigma(w_{r0}x + b_{r0}) + \sum_{r=1}^m a_{r0} (w_r x + b_r) \sigma(w_{r0}x + b_{r0}) + \Delta_{np}^T \\ &\stackrel{(i)}{\leq} 16c_1 (c_6 + c_7) \epsilon_a \log m + m \left(2c_1 \epsilon_a \sqrt{2 \log m} \right) (2\eta\bar{\Delta}T) + \Delta_{np}^T \\ &\leq 16c_1 (c_6 + c_7) \epsilon_a \log m + m \left(2c_1 \epsilon_a \sqrt{2 \log m} \right) (2\eta\bar{\Delta}T) + \Delta_{np}^T \\ &\stackrel{(ii)}{\leq} 16c_1 (c_6 + c_7) \epsilon_a \log m + m \left(48c_1^2 \epsilon_a^2 \log m \left(\frac{U_h^2 \pi}{4m\epsilon_a^2 \epsilon} \right) \right) + \Delta_{np}^T \\ &\leq O\left(\frac{U_h^2 \log m}{\epsilon} \right) \end{aligned}$$

where inequality (i) uses eq. (12), lemma 15 and eq.(18). The inequality (ii) uses our choices of η and T from eq.(30) and lower bound on m from eq.(31). Define K as upper bound on $\sup_{N \in \mathcal{F}, x \in [-1, 1]} N(x)$.

$$K := O\left(\frac{U_h^2 \log m}{\epsilon} \right) \quad (37)$$

Using upper bound on $\sup_{N \in \mathcal{F}, x \in [-1, 1]} N(x)$ and eq.(36), we get upper bound on $\hat{L}(b)$.

$$b = K + K + \hat{L}(0, 0, \dots, 0, 0) \leq 2K + 2$$

Using value of b in eq.(34) and lemma 11, with at least $1 - \delta - \frac{1}{c_1}$ probability, we get

$$\begin{aligned} \sup_{N \in \mathcal{F}} \left| \mathbb{E}_{x \in \mathcal{D}} \left[\hat{L}(N(\tau_1(x)), \dots, N(\tau_{\tilde{n}}(x)), N(x)) \right] - \frac{1}{n} \sum_{i=1}^n \hat{L}(N(\tau_1(x_i)), \dots, N(\tau_{\tilde{n}}(x_i)), N(x_i)) \right| \\ \leq 4\sqrt{2}(Q+1) \frac{8c_1 \epsilon_a \eta \bar{\Delta} T m \sqrt{2 \log m}}{\sqrt{n}} + (2K + 2) \sqrt{\frac{\log \frac{1}{\delta}}{2n}} \end{aligned}$$

We use $\delta = 0.01$ and choose n which satisfies following condition.

$$n \geq O\left(\frac{(M_{\hat{L}} - m_{\hat{L}})^2 (Q+1)^2 U_h^4 (\log m)^2}{\epsilon^2}\right) \quad (38)$$

Using above n , with at least 0.98 probability, we get

$$\sup_{N \in \mathcal{F}} \left| \mathbb{E}_{x \in \mathcal{D}} \left[\hat{L}(N(\tau_1(x)), \dots, N(\tau_{\tilde{n}}(x)), N(x)) \right] - \frac{1}{n} \sum_{i=1}^n \hat{L}(N(\tau_1(x_i)), \dots, N(\tau_{\tilde{n}}(x_i)), N(x_i)) \right| \leq \epsilon$$

□

Lemma 13. (Concentration on approximated loss of target function) Suppose n is sufficiently high such that it satisfies following condition.

$$n \geq O\left(\frac{(M_{\hat{L}} - m_{\hat{L}})^2 (Q+1)^2 U_h^4 (\log m)^2}{\epsilon^2}\right)$$

If n satisfies above condition, then with at least 0.9999 probability, population loss of target function $F^{*'}$ is close to empirical loss i.e.

$$\left| \mathbb{E}_{x \sim \mathcal{D}} \left[\hat{L}(F^{*'}, x) \right] - \hat{L}(F^{*'}, \mathcal{X}) \right| \leq \epsilon$$

Proof. Finding minimum value ($m_{\hat{L}}$) and maximum value ($M_{\hat{L}}$) of loss function \hat{L} ,

$$\begin{aligned} \sup_x \hat{L}(F^{*'}, x) &= \sup_x \sum_{i=1}^{\tilde{n}} \tilde{w}_i F^{*'}(\tau_i(x)) - \log(F^{*'}(x)) \leq 2M_{F^{*'}} - \log(m_{F^{*'}}) = M_{\hat{L}} \\ \inf_x \hat{L}(F^{*'}, x) &= \inf_x \sum_{i=1}^{\tilde{n}} \tilde{w}_i F^{*'}(\tau_i(x)) - \log(F^{*'}(x)) \leq 2m_{F^{*'}} - \log(M_{F^{*'}}) = m_{\hat{L}} \end{aligned}$$

where $M_{F^{*'}} = \max_{x \in [-1, 1]} F^{*'}(x)$ and $m_{F^{*'}} = \min_{x \in [-1, 1]} F^{*'}(x)$. Using Hoeffding's inequality,

$$\Pr\left(\left| \mathbb{E}_{x \sim \mathcal{D}} \left[\hat{L}(F^{*'}, \mathcal{X}) \right] - \hat{L}(F^{*'}, \mathcal{X}) \right| \geq \epsilon\right) \leq \exp\left(-\frac{2n\epsilon^2}{(M_{\hat{L}} - m_{\hat{L}})^2}\right)$$

Taking n as

$$n \geq O\left(\frac{(M_{\hat{L}} - m_{\hat{L}})^2 (Q+1)^2 U_h^4 (\log m)^2}{\epsilon^2}\right)$$

With at least probability $1 - \exp\left(-\frac{2n\epsilon^2}{(M_{\hat{L}} - m_{\hat{L}})^2}\right)$,

$$\left| \mathbb{E}_{x \sim \mathcal{D}} \left[\hat{L}(F^{*'}, x) \right] - \hat{L}(F^{*'}, \mathcal{X}) \right| \leq \epsilon \quad (39)$$

□

Corollary 2. Under same setting as lemma 10 and

$$n \geq O\left(\frac{(M_{\hat{L}} - m_{\hat{L}})^2 (Q+1)^2 U_h^4 (\log m)^2}{\epsilon^2}\right)$$

then with at least $0.92 - 2 \exp\left(-\frac{m}{8\epsilon_a^2}\right) - \exp\left(-\frac{\epsilon^2 m}{128c_1^2 U_h^2 \log m}\right) - \exp\left(-\frac{2n\epsilon^2}{(M_{\hat{L}} - m_{\hat{L}})^2}\right)$ probability, we get

$$\mathbb{E}_{sgd} \left[\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_{x \sim \mathcal{D}} \left[\hat{L}(f'_t, x) \right] \right] - \mathbb{E}_{x \sim \mathcal{D}} \left[\hat{L}(F^{*'}, x) \right] \leq O(\epsilon)$$

Proof. Using lemma 10, lemma 12 and lemma 13, with at least $0.92 - 2 \exp\left(-\frac{m}{8\epsilon_a^2}\right) - \exp\left(-\frac{\epsilon^2 m}{128c_1^2 U_h^2 \log m}\right) - \exp\left(-\frac{2n\epsilon^2}{(M_{\hat{L}} - m_{\hat{L}})^2}\right)$ probability, we get

$$\mathbb{E}_{sgd} \left[\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_{x \sim \mathcal{D}} \left[\hat{L}(f'_t, x) \right] \right] - \mathbb{E}_{x \sim \mathcal{D}} \left[\hat{L}(F^{*'}, x) \right] \leq O(\epsilon)$$

□

Theorem 2. (loss function is close to optimal) Taking same choice of parameters as in lemma 10 and $Q \geq \frac{4M_{F^{*''}} + 4K_2}{\epsilon}$, with at least $0.92 - 2 \exp\left(-\frac{m}{8\epsilon_a^2}\right) - \exp\left(-\frac{\epsilon^2 m}{128c_1^2 U_h^2 \log m}\right) - \exp\left(-\frac{2n\epsilon^2}{(M_{\hat{L}} - m_{\hat{L}})^2}\right)$ probability, we have

$$\mathbb{E}_{sgd} \left[\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_{x \sim \mathcal{D}} [L(f_t, x)] \right] - \mathbb{E}_{x \sim \mathcal{D}} [L(F^*, x)] \leq O(\epsilon)$$

where

$$\begin{aligned} M_{F^{*''}} &= \sup_x F^{*''}(x) \\ M_{\hat{L}} &= 2M_{F^{*'}} - \log(m_{F^{*'}}) \\ m_{\hat{L}} &= 2m_{F^{*'}} - \log(M_{F^{*'}}) \\ M_{F^{*'}} &= \sup_x F^{*'}(x) \\ m_{F^{*'}} &= \inf_x F^{*'}(x) \end{aligned}$$

Proof. First, we will try to bound

$$\begin{aligned} \left| \hat{L}(F^{*'}, x) - L(F^*, x) \right| &\leq \left| \sum_{i=1}^Q \Delta_x F^{*'}(\tau_i(x)) - F^*(x) \right| \\ &\leq \frac{2M_{F^{*''}}}{Q} \end{aligned}$$

Similarly, bounding error for f'_t , we will get

$$\begin{aligned} \left| \hat{L}(f'_t, x) - L(f_t, x) \right| &\leq \left| \sum_{i=1}^Q \Delta_x f'_t(\tau_i(x)) - f_t(x) \right| \\ &\leq \frac{2(\sup_x f''_t(x))}{Q} \end{aligned}$$

To get $\sup_x f_t''(x)$, we will use eq.(37).

$$\begin{aligned}
\sup_x f_t''(x) &\leq \sup_x |N_t'(x)| \\
&\leq \sup_x \sum_{r=1}^m a_{r0} \sigma'((w_{r0} + w_r^t)x + b_{r0} + b_r^t) (w_{r0} + w_r^t) \\
&\leq \sup_x \sum_{r=1}^m a_{r0} (w_{r0} + w_r^t) \mathbb{I}[(w_{r0} + w_r^t)x + b_{r0} + b_r^t] \\
&\leq \sup_x \sum_{r \in \mathcal{H}} a_{r0} (w_{r0} + w_r^t) \mathbb{I}[w_{r0}x + b_{r0}] + \sum_{r \notin \mathcal{H}} a_{r0} (w_{r0} + w_r^t) \mathbb{I}[(w_{r0} + w_r^t)x + b_{r0} + b_r^t] \\
&\leq \sup_x \sum_{r \in \mathcal{H}} a_{r0} w_{r0} \mathbb{I}[w_{r0}x + b_{r0}] + \sum_{r \in \mathcal{H}} a_{r0} w_r^t \mathbb{I}[w_{r0}x + b_{r0}] \\
&\quad + \sum_{r \notin \mathcal{H}} a_{r0} w_{r0} \mathbb{I}[(w_{r0} + w_r^t)x + b_{r0} + b_r^t] + \sum_{r \notin \mathcal{H}} a_{r0} w_r^t \mathbb{I}[(w_{r0} + w_r^t)x + b_{r0} + b_r^t] \\
&\stackrel{(i)}{\leq} 16c_1 c_6 \epsilon_a (\log m) + m \left(2c_1 \epsilon_a \sqrt{2 \log m} \right) (\eta \bar{\Delta} T) \\
&\quad + \left(c_2 \frac{4\sqrt{2}\eta m \sqrt{m} \bar{\Delta} t}{\sqrt{\pi}} \right) \left(2c_1 \epsilon_a \sqrt{2 \log m} \right) \left(\frac{2c_6 \sqrt{2 \log m}}{\sqrt{m}} \right) \\
&\quad + \left(c_2 \frac{4\sqrt{2}\eta m \sqrt{m} \bar{\Delta} t}{\sqrt{\pi}} \right) \left(2c_1 \epsilon_a \sqrt{2 \log m} \right) (\eta \bar{\Delta} T) \\
&\leq O(\epsilon) + m \left(2c_1 \epsilon_a \sqrt{2 \log m} \right) \left(\bar{\Delta} \frac{U_h^2 \pi}{4m \epsilon_a^2 \epsilon} \right) \\
&\quad + \left(c_2 \frac{4\sqrt{2}m \sqrt{m} \bar{\Delta}}{\sqrt{\pi}} \right) \left(2c_1 \epsilon_a \sqrt{2 \log m} \right) \left(\frac{2c_6 \sqrt{2 \log m}}{\sqrt{m}} \right) \frac{U_h^2 \pi}{4m \epsilon_a^2 \epsilon} \\
&\quad + \left(c_2 \frac{4\sqrt{2}m \sqrt{m} \bar{\Delta}}{\sqrt{\pi}} \right) \left(2c_1 \epsilon_a \sqrt{2 \log m} \right) \left(\bar{\Delta} \left(\frac{U_h^2 \pi}{4m \epsilon_a^2 \epsilon} \right)^2 \right) \\
&\leq O(\epsilon) + O\left(\frac{U_h^2 \log m}{\epsilon}\right) + O\left(\frac{U_h^2 (\log m)^{1.5}}{\epsilon}\right) + O(\epsilon) \\
&\leq O\left(\frac{U_h^2 (\log m)^{1.5}}{\epsilon}\right)
\end{aligned}$$

Define K_2 as upper bound on $\sup_x f_t''(x)$,

$$K_2 = O\left(\frac{U_h^2 (\log m)^{1.5}}{\epsilon}\right)$$

Taking Q as

$$Q \geq \frac{2M_{F^{**}} + 2K_2}{\epsilon} \quad (40)$$

Using given value of Q , we get that

$$\left| \hat{L}(F^{*'}, x) - L(F^*, x) \right| \leq \epsilon \quad (41)$$

$$\left| \hat{L}(f_t', x) - L(f_t, x) \right| \leq \epsilon \quad (42)$$

Using these relations, we get

$$\mathbb{E}_{sgd} \left[\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}_{x \sim \mathcal{D}} [L(f_t, x)] \right] - \mathbb{E}_{x \sim \mathcal{D}} [L(F^*, x)] \leq O(\epsilon)$$

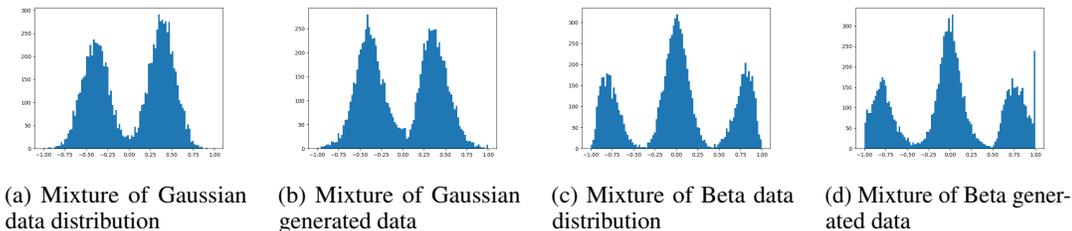


Figure 2: Comparison of data distribution and generated data for mixture of Gaussian and beta distributions

By definition of KL divergence, we get

$$\mathbb{E}_{sgd} \left[\frac{1}{T} \sum_{t=0}^{T-1} \text{KL}(p_{F^*,Z} || p_{f_t,Z}) \right] \leq O(\epsilon)$$

□

E ADDITIONAL EXPERIMENTS

In this section, we show experimental results on synthetic 1D data to support our theoretical findings. We use the same architecture and initialization as described in Sec. 2 for both constrained and unconstrained normalizing flows. In all our experiments, we fix the weights of the output layer and train the weights and biases of the hidden layer. For training, we use mini-batch SGD with batch size 32. We use 2 datasets each with 10,000 data points. One of the dataset is mixture of 2 Gaussians and other one is mixture of 3 beta distributions. All results are averaged over 3 different iterations.

E.1 RESULTS FOR UNCONSTRAINED NORMALIZING FLOW

In Fig. 2, we show comparison of data distribution and generated data distribution for unconstrained normalizing flow. Unconstrained normalizing flow with exponential target distribution learns the data distribution well. We study the effect of overparameterization on L_2 -norm of W^t and B^t and convergence speed. To reproduce situation similar to the theoretical analyses for unconstrained normalizing flow, we choose learning rate as $\frac{c}{m}$ where c is a constant. The first row of Figure 3 contains results for mixture of Gaussians dataset and the second row contains results for mixtures of beta distributions dataset. From Fig. 3, we see that L_2 -norm of W^t and B^t decreases with increasing m . Moreover, the change is proportional to $1/\sqrt{m}$ which is similar to the bound in theoretical result. From the last column of Fig. 3, we see that the training speed for different values of m remains almost constant. Our choice of T in theoretical analysis also poly-logarithmically depends upon m . We obtained similar results for Gaussian distribution as well.

E.2 RESULTS FOR CONSTRAINED NORMALIZING FLOW

In Sec. 2.1, we suggested that high overparameterization may adversely affect training for constrained normalizing flows. We now give experimental evidence for this in Figs. 4, 5 and 6. In Fig.4, we see that for neural network with $m = 100$ and $m = 400$, the training loss decreases stably. In Figs. 5 and 6, we see that as we increase the learning rate, training becomes more stable for larger m . Note that for learning rate equal to 0.025 and 0.0125, constrained normalizing flow with $m = 1600$ doesn't learn anything due to small learning rate. We observe that the L_2 -norms of W^t and B^t for $m = 6400$ are at least as large as those of $m = 1600$. On both datasets, as we increase the learning rate, L_2 -norm of B^t increases (except for learning rate=0.05 of mixture of beta distribution) and learning of constrained normalizing flow becomes more and more stable. These observations support our claim in Sec.2.1 that for learning and approximation of overparameterized constrained normalizing flow, neural networks need large L_2 -norms of W^t and B^t .

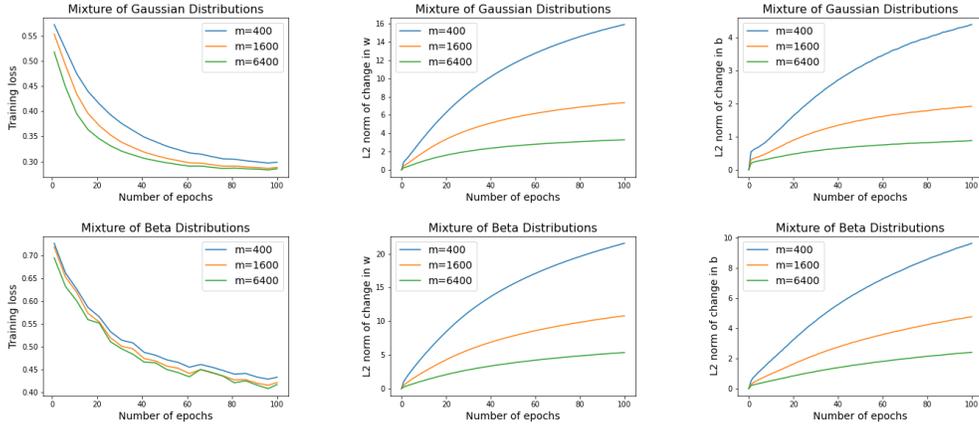


Figure 3: Effect of over-parameterization on training of unconstrained normalizing flow on mixture of Gaussian and mixture of beta distributions

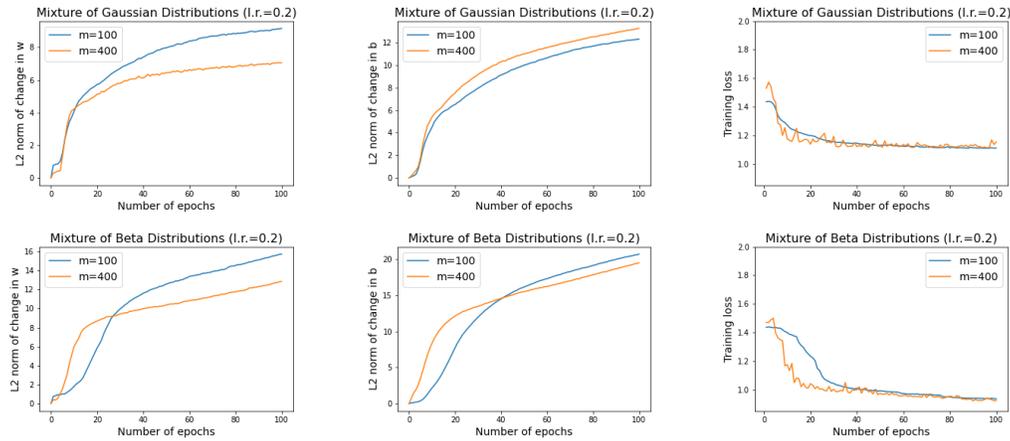


Figure 4: Effect of over-parameterization on training of small sized constrained normalizing flow

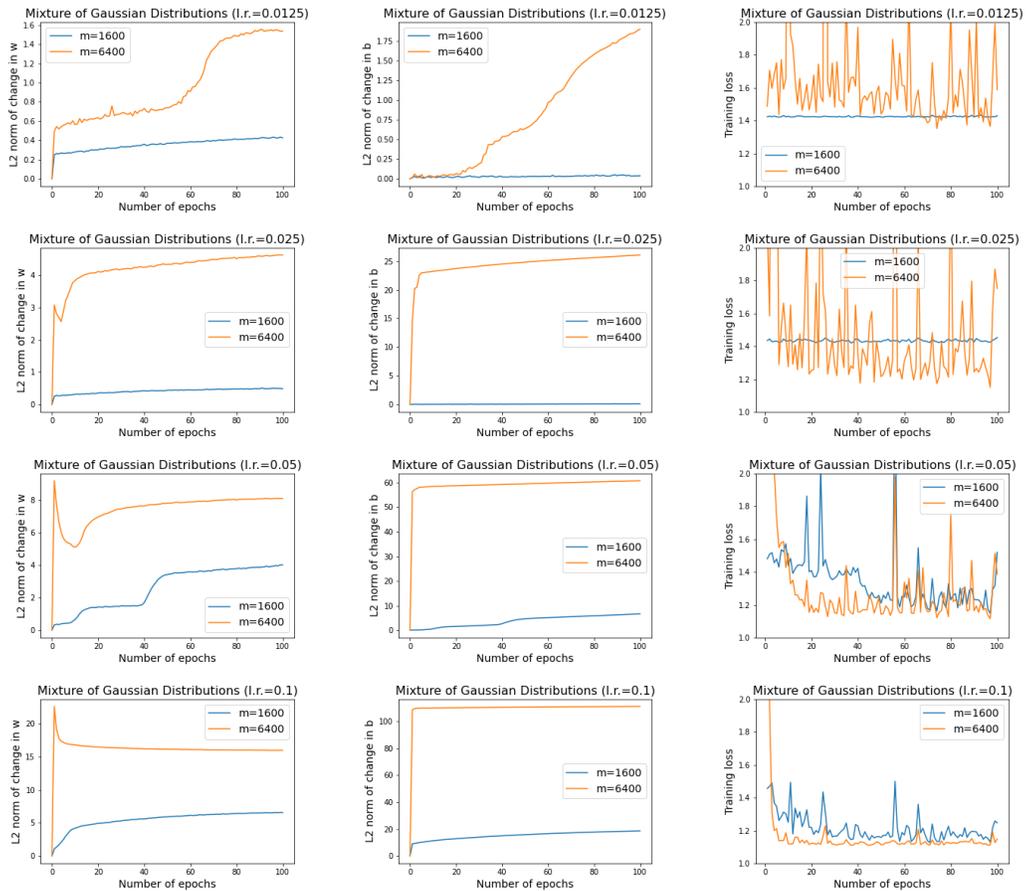


Figure 5: Effect of over-parameterization on training of constrained normalizing flow on mixture of Gaussian dataset for number of hidden layers $m = 1600, 6400$

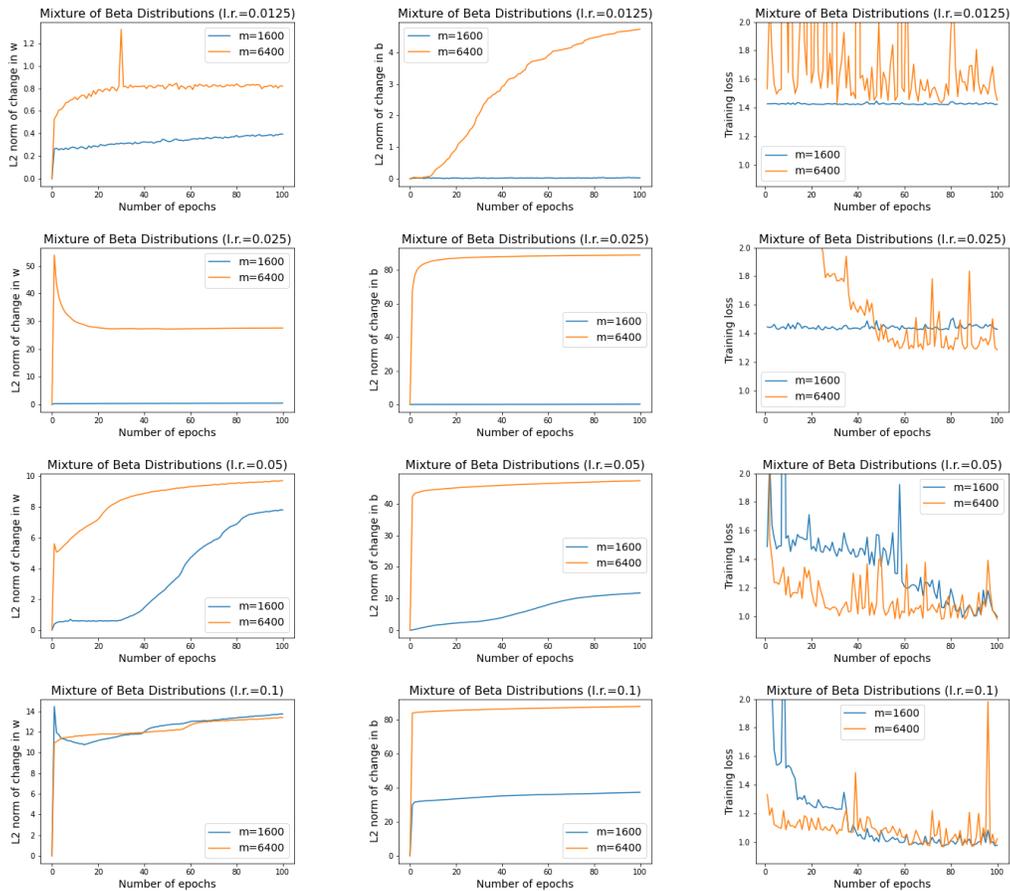


Figure 6: Effect of over-parameterization on training of constrained normalizing flow on mixture of beta distribution dataset for number of hidden layers $m = 1600, 6400$

F USEFUL FACTS

Lemma 14. Suppose $Z_k \sim \mathcal{N}(0, \sigma^2)$ and $Y = \sum_{k=1}^n Z_k^2$ is chi-squared distribution with following property for all $t \in (0, 1)$.

$$\Pr \left[\left| \frac{1}{n} \sum_{k=1}^n Z_k^2 - \sigma^2 \right| \geq t \right] \leq 2 \exp \left(-\frac{nt^2}{8\sigma^4} \right)$$

Proof. From example 2.11 from Wainwright (2019), for $Z'_k \sim \mathcal{N}(0, 1)$ and $Y = \sum_{k=1}^n Z_k'^2$ is chi-squared distribution with following property for all $t \in (0, 1)$.

$$\Pr \left[\left| \frac{1}{n} \sum_{k=1}^n Z_k'^2 - 1 \right| \geq t \right] \leq 2 \exp \left(-\frac{nt^2}{8} \right)$$

Using above equation for $\frac{Z_k}{\sigma}$,

$$\Pr \left[\left| \frac{1}{n} \sum_{k=1}^n \frac{Z_k^2}{\sigma^2} - 1 \right| \geq \frac{t}{\sigma^2} \right] \leq 2 \exp \left(-\frac{nt^2}{8\sigma^4} \right)$$

$$\Pr \left[\left| \frac{1}{n} \sum_{k=1}^n Z_k^2 - \sigma^2 \right| \geq t \right] \leq 2 \exp \left(-\frac{nt^2}{8\sigma^4} \right)$$

□

Lemma 15. Let X_1, X_2, \dots, X_n be independent random variables from $\mathcal{N}(0, \sigma^2)$, then with at least $1 - \frac{1}{c_1}$ probability, following holds.

$$\max_{i \in \{1, 2, \dots, n\}} |X_i| \leq 2c_1 \sigma \sqrt{2 \log n}$$

Proof. From Romberg (2012),

$$\mathbb{E} \left[\max_{i \in \{1, 2, \dots, n\}} |X_i| \right] \leq \sigma \left(\sqrt{2 \log n} + 1 \right) \leq 2\sigma \left(\sqrt{2 \log n} \right)$$

Assuming $n \geq 2$, the last inequality follows. Using Markov's inequality,

$$\Pr \left(\max_{i \in \{1, 2, \dots, n\}} |X_i| \geq 2c_1 \sigma \left(\sqrt{2 \log n} \right) \right) \leq \frac{1}{c_1}$$

$$\Pr \left(\max_{i \in \{1, 2, \dots, n\}} |X_i| \leq 2c_1 \sigma \left(\sqrt{2 \log n} \right) \right) \geq 1 - \frac{1}{c_1}$$

s

□

Lemma 16. For standard Gaussian random variable X from $\mathcal{N}(0, \sigma^2)$, following anti-concentration inequality holds.

$$\Pr(|X| \leq R) \leq \frac{2R}{\sigma\sqrt{2\pi}}$$

Proof. (From Du et al. (2018)) For standard Gaussian random variable $\frac{X}{\sigma}$,

$$\Pr \left(\left| \frac{X}{\sigma} \right| \leq R \right) \leq \frac{2R}{\sqrt{2\pi}}$$

Using $R = \frac{R'}{\sigma}$, we get the required result.

□

Lemma 17. Suppose function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ is L_g -Lipschitz continuous and L_i -coordinate wise Lipschitz continuous i.e.

$$\begin{aligned} |f(\mathbf{a}) - f(\mathbf{b})| &\leq L_g \|\mathbf{a} - \mathbf{b}\| \\ &\quad \forall \mathbf{a}, \mathbf{b} \in \mathbb{R}^d \quad (\text{Standard Lipschitz continuity}) \\ |f(a_1, a_2, \dots, a_i, \dots, a_d) - f(a_1, a_2, \dots, b_i, \dots, a_d)| &\leq L_i |a_i - b_i| \\ &\quad \forall a_1, a_2, \dots, a_i, \dots, a_d, b_i \in \mathbb{R} \text{ and } \forall i \in [d] \quad (\text{Coordinate-wise Lipschitz continuity}) \end{aligned}$$

If a function f satisfies L_i -coordinate wise Lipschitz continuity for all i , then function f follows following inequality.

$$|f(a_1, a_2, \dots, a_d) - f(b_1, b_2, \dots, b_d)| \leq \sum_{i=1}^d L_i |a_i - b_i|$$

Moreover, the function f also satisfies standard Lipschitz continuity with L_g Lipschitz constant where inequality between L_g and L_i is as follows.

$$L_g \leq \sqrt{\sum_{i=1}^d L_i^2}$$

Proof. Define $\mathbf{a} = (a_1, a_2, \dots, a_d)$ and $\mathbf{b} = (b_1, b_2, \dots, b_d)$.

$$\begin{aligned} |f(a_1, a_2, \dots, a_d) - f(b_1, b_2, \dots, b_d)| &\leq |f(a_1, a_2, \dots, a_d) - f(b_1, a_2, \dots, a_d)| \\ &\quad + |f(b_1, a_2, a_3, \dots, a_d) - f(b_1, b_2, a_3, \dots, a_d)| \\ &\quad + |f(b_1, b_2, a_3, \dots, a_d) - f(b_1, b_2, b_3, \dots, a_d)| \\ &\quad + \dots + |f(b_1, b_2, \dots, b_{d-1}, a_d) - f(b_1, b_2, b_3, \dots, b_d)| \\ &\leq L_1 |a_1 - b_1| + L_2 |a_2 - b_2| + \dots + L_d |a_d - b_d| \\ &\leq \sqrt{\sum_{i=1}^d L_i^2} \|\mathbf{a} - \mathbf{b}\|_2 \end{aligned}$$

where last inequality follows from Cauchy-Schwarz inequality. \square

Fact 1. (Hoeffding's inequality on Binomial random variable) If we have a binomial random variable with parameters n (total number of trials) and p (probability of success). For number successful trial $k \leq np$, following inequality holds.

$$\Pr(X \leq k) \leq \exp\left(-2n \left(p - \frac{k}{n}\right)^2\right)$$

Fact 2. For a gaussian random variable $X \sim \mathcal{N}(0, \sigma^2)$, $\forall t \in (0, \sigma)$, we have

$$\Pr(|X| \geq t) \geq 1 - \frac{4t}{5\sigma}$$

Fact 3. The sum of reciprocals of the squares of the natural numbers is given by

$$\sum_{n=1}^{\infty} \frac{1}{n^2} = \frac{\pi^2}{6} \leq 2$$

Fact 4. (Theorem 3.1(r'_5) of Li & Yeh (2013)) For any $\alpha > 1$ and $x \in \left[0, \frac{1}{\alpha-1}\right)$,

$$(1+x)^\alpha \leq \frac{1}{1 - \frac{\alpha x}{1+x}} = 1 + \frac{\alpha x}{1 - (\alpha-1)x}$$

Fact 5. *If Arithmetic-Geometric Progression(AGP) is as follows.*

$$a, (a + d)r, (a + 2d)r^2, (a + 3d)r^3, \dots, [a + (n - 1)d] r^{n-1}$$

where a is the initial term, d is the common difference and r is the common ratio. The sum of the first n terms of the AGP (S_n) is given by

$$S_n = \frac{a - [a + (n - 1)d] r^n}{1 - r} + \frac{dr(1 - r^{n-1})}{(1 - r)^2}$$

Fact 6. (Lemma A.3 from Ji et al. (2020)) *The Fourier transform of f is defined as*

$$\hat{f}(w) = \int f(x)e^{2\pi iwx} dx$$

The polar decomposition of the Fourier transform \hat{f} is $\hat{f}(w) = |\hat{f}(w)| e^{2\pi i\theta_f(w)}$ with $|\theta_f(w)| \leq 1$. The Fourier transform \hat{f} follows below properties.

1. $|\hat{f}(w)| \leq \|f\|_{L_1}$ for any real number w .
2. Let $\alpha > 0$ be given and define $\beta := \frac{1}{2\pi\alpha}$. G_α is Gaussian with coordinate-wise variance α^2 . Then $|\hat{G}_\alpha| = \hat{G}_\alpha$ (meaning \hat{G}_α has no radial component) and

$$\hat{G}_\alpha(w) = \frac{1}{\sqrt{2\pi\alpha^2}} G_\beta(w) = \sqrt{2\pi\beta^2} G_\beta(w) = \sqrt{2\pi} G(w/\beta)$$

Definition 1. Let \mathcal{F} be a set of functions $\mathbb{R}^d \rightarrow \mathbb{R}$ and $\mathcal{X} = (x_1, x_2, \dots, x_n)$ be a finite set of samples. The empirical Rademacher complexity of \mathcal{F} with respect to \mathcal{X} is defined by

$$\hat{\mathcal{R}}(\mathcal{X}; \mathcal{F}) = \mathbb{E}_{\xi \sim \{\pm 1\}^n} \left[\sup_{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^n \xi_i f(x_i) \right]$$

Proposition 1. *The Rademacher complexity have following properties.*

- a. Suppose $|x| \leq 1$ for all \mathcal{X} . The class $\mathcal{F} = \{x \rightarrow wx + b \mid |w| \leq B, |b| \leq B\}$ has Rademacher complexity $\hat{\mathcal{R}}(\mathcal{X}, \mathcal{F}) \leq \frac{2B}{\sqrt{n}}$
- b. Given $\mathcal{F}_1, \mathcal{F}_2$ classes of functions, then $\hat{\mathcal{R}}(\mathcal{X}; \mathcal{F}_1 + \mathcal{F}_2) = \hat{\mathcal{R}}(\mathcal{X}; \mathcal{F}_1) + \hat{\mathcal{R}}(\mathcal{X}; \mathcal{F}_2)$
- c. Given $\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_m$ classes of functions from $\mathcal{X} \rightarrow \mathbb{R}$ and suppose $w \in \mathbb{R}^m$ is a fixed vector; then $\mathcal{F}' = \{x \rightarrow \sum_{r=1}^m w_r \sigma(f_r(x)) \mid f_r \in \mathcal{F}_r\}$ satisfies $\hat{\mathcal{R}}(\mathcal{X}; \mathcal{F}') \leq 2\|w\|_1 \max_{r \in [m]} \hat{\mathcal{R}}(\mathcal{X}; \mathcal{F}_r)$ where σ is 1-Lipschitz continuous function.

Proof. The b and c parts of the proposition are from Allen-Zhu et al. (2019) . Proof of the a part is as following.

$$\begin{aligned} \hat{\mathcal{R}}(\mathcal{X}, \mathcal{F}) &= \frac{1}{n} \mathbb{E}_{\xi \sim \{\pm 1\}^n} \left[\sup_{f \in \mathcal{F}} w \left(\sum_{i=1}^n \xi_i x_i \right) + b \left(\sum_{i=1}^n \xi_i \right) \right] \\ &= \frac{B}{n} \mathbb{E}_{\xi \sim \{\pm 1\}^n} \left[\left| \sum_{i=1}^n \xi_i x_i \right| + \left| \sum_{i=1}^n \xi_i \right| \right] \end{aligned}$$

Using Jensen's inequality, we get

$$\mathbb{E}_{\xi \sim \{\pm 1\}^n} \left[\left| \sum_{i=1}^n \xi_i x_i \right| \right] = \mathbb{E}_\xi \left[\left(\left| \sum_{i=1}^n \xi_i x_i \right|^2 \right)^{0.5} \right] \leq \left(\mathbb{E}_\xi \left[\left| \sum_{i=1}^n \xi_i x_i \right|^2 \right] \right)^{0.5}$$

Using independence of ξ_i for all $i \in [n]$, we get

$$\begin{aligned}\mathbb{E}_\xi \left[\left(\sum_{i=1}^n \xi_i x_i \right)^2 \right] &= \mathbb{E}_\xi \left[\sum_{i,j} \xi_i \xi_j x_i x_j \right] \\ &= \sum_{i=1}^n |x_i|^2 \leq n\end{aligned}$$

Using same technique, we will get same bound for $\mathbb{E}_\xi [|\sum_{i=1}^n \xi_i|]$.

$$\hat{\mathcal{R}}(\mathcal{X}, \mathcal{F}) \leq \frac{2B}{\sqrt{n}}$$

□

Fact 7. (Rademacher Complexity) If $\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_k$ are k classes of functions $\mathbb{R}^d \rightarrow \mathbb{R}$ and $L_x : \mathbb{R}^d \rightarrow [-b, b]$ is L_g -Lipschitz continuous function for any $x \sim \mathcal{D}$, then

$$\sup_{f_1 \in \mathcal{F}_1, \dots, f_k \in \mathcal{F}_k} \left| \mathbb{E}_{x \in \mathcal{D}} [L_x(f_1(x), \dots, f_k(x))] - \frac{1}{n} \sum_{i=1}^n L_x(f_1(x_i), \dots, f_k(x_i)) \right| \leq 2\hat{\mathcal{R}}(\mathcal{X}; \mathcal{L}) + b\sqrt{\frac{\log \frac{1}{\delta}}{2n}}$$

where \mathcal{L} is set of all functions L_x . Using vector contraction inequality from Maurer (2016), we get

$$\begin{aligned}\sup_{f_1 \in \mathcal{F}_1, \dots, f_k \in \mathcal{F}_k} \left| \mathbb{E}_{x \in \mathcal{D}} [L_x(f_1(x), \dots, f_k(x))] - \frac{1}{n} \sum_{i=1}^n L_x(f_1(x_i), \dots, f_k(x_i)) \right| \\ \leq 2\sqrt{2}L_g \left(\sum_{i=1}^k \hat{\mathcal{R}}(\mathcal{X}; \mathcal{F}_i) \right) + b\sqrt{\frac{\log \frac{1}{\delta}}{2n}}\end{aligned}$$